NASA/CR-1999-209728

Semiannual Report

April 1, 1999 through September 30, 1999

Institute for Computer Applications in Science and Engineering
NASA Langley Research Center
Hampton, VA

Operated by Universities Space Research Association

National Aeronautics and
Space Administration

Langley Research Center
Hampton, Virginia 23681-2199

Prepared for Langley Research Center
under Contract NAS1-97046

November 1999

Available from the following:

NASA Center for AeroSpace Information (CASI)
7121 Standard Drive
Hanover, MD 21076-1320
(301) 621-0390

National Technical Information Service (NTIS)
5285 Port Royal Road
Springfield, VA 22161-2171
(703) 487-4650

CONTENTS

Page

Introduction ii

Research in Progress

Applied and Numerical Mathematics 1

Physical Sciences 20

Computer Science 27

Reports and Abstracts 38

ICASE Interim Reports 46

ICASE Colloquia 47

ICASE Summer Activities 51

Other Activities 55

ICASE Staff 56

INTRODUCTION

The Institute for Computer Applications in Science and Engineering (ICASE)* is operated at the Langley Research Center (LaRC) of NASA by the Universities Space Research Association (USRA) under a contract with the Center. USRA is a nonprofit consortium of major U.S. colleges and universities.

The Institute conducts unclassified basic research in applied mathematics, numerical analysis and algorithm development, fluid mechanics, and computer science in order to extend and improve problem-solving capabilities in science and engineering, particularly in the areas of aeronautics and space research.

ICASE has a small permanent staff. Research is conducted primarily by visiting scientists from universities and industry who have resident appointments for limited periods of time as well as by visiting and resident consultants. Members of NASA's research staff may also be residents at ICASE for limited periods.

The major categories of the current ICASE research program are:

- Applied and numerical mathematics, including multidisciplinary design optimization;
- Theoretical and computational research in fluid mechanics in selected areas of interest to LaRC, such as transition, turbulence, flow control, and acoustics; and
- Applied computer science: system software, systems engineering, and parallel algorithms.

ICASE reports are primarily considered to be preprints of manuscripts that have been submitted to appropriate research journals or that are to appear in conference proceedings. A list of these reports for the period April 1, 1999 through September 30, 1999 is given in the Reports and Abstracts section which follows a brief description of the research in progress.

*ICASE is operated at NASA Langley Research Center, Hampton, VA, under the National Aeronautics and Space Administration, NASA Contract No. NAS1-97046. Financial support was provided by NASA Contract Nos. NAS1-97046, NAS1-19480, NAS1-18605, NAS1-18107, NAS1-17070, NAS1-17130, NAS1-15810, NAS1-16394, NAS1-14101, and NAS1-14472.

RESEARCH IN PROGRESS

APPLIED AND NUMERICAL MATHEMATICS

Absorbing Layers for Far-field Boundary Conditions

Saul Abarbanel

The ability to accurately simulate wave phenomena is important in several physical fields, e.g., electromagnetics, ambient acoustics, advected acoustics associated with a mean flow, elasticity, and seismology. Often, due to limited computing resources, the numerical simulations of such problems must be confined to truncated domains much smaller than the physical space over which the wave phenomena take place. In such cases, spurious numerical reflections of outgoing waves from the boundaries of the numerical domain can falsify the computational results. This is particularly troublesome in cases where higher order of accuracy is required by mode resolution, storage availability, etc. The absorbing layer approach deals with this type of problem.

In previous work it was shown how, in the two-dimensional case, posing a traveling-plane-wave solution allows one to construct the modified partial differential equations in the absorbing layers. Applications were made to problems in advected acoustics as well as electromagnetics. In this reporting period the effort was directed towards extending the work to the three-dimensional case. We were able to construct such three-dimensional layers in the electromagnetic case. We believe our formulation will turn out to be at least competitive with that based on the Lorentz material approach.

This work is part of ongoing research in collaboration with David Gottlieb and Jan Hesthaven (Brown University).

Closed-loop Separation Control Using Oscillatory Flow Excitation

Brian G. Allan

Experimental investigations have demonstrated that the introduction of periodic vortical excitations, slightly upstream of a separated boundary layer, can effectively delay boundary layer separation. Recent experiments by Seifert have demonstrated flow separation control over a hump model, at high Reynolds numbers, using the 0.3-meter Transonic Cryogenic Tunnel at NASA Langley. The degree of separation behind the hump can be quantified by dC_P/dx, a pressure gradient downstream of the hump. Since the degree of separation is proportional to the aerodynamic lift generated by the hump model, the lift can be controlled by varying the amplitude of the jet. The objective of this investigation is to incorporate feedback control into the experimental hump system where the closed-loop system will track a given dC_P/dx command.

A feedback controller was developed using a simple dynamic model of the flow system and was designed to track a desired dC_P/dx input command. Wind tunnel experiments using the hump model were conducted at the 0.3-meter Transonic Cryogenic Tunnel at NASA Langley. The wind tunnel tests were conducted at a Reynolds number of 16 million and a Mach number of 0.25. Open-loop experiments were used to model the dynamics of the flow system which were then used to design the feedback controller. The closed-loop experiments demonstrated good tracking and improved transient characteristics as compared to the open-loop results.

Future work will involve interactions with other flow control experiments currently being conducted at NASA Langley. These interactions will incorporate the development of control designs for closed-loop experimental tests of novel flow control devices. This research was conducted in collaboration with Jer-Nan Juang (NASA Langley), David Raney (NASA Langley), and Avi Seifert (NRC).

Approximations of the Newton Step for Large Scale Optimization Problems

Eyal Arian

Quasi-Newton methods for large scale optimization problems are powerful but suffer an initial slow convergence rate. Our goal is to develop a new iterative method, for the solution of large scale optimization problems, that will allow a better approximation for the Newton step right from the first optimization steps.

In the course of the optimization process, systems of linear equations are constructed that contain the linearized state operator and its adjoint. These have to be solved at each iteration to achieve convergence of the iterates to the Newton step. We are investigating a defect-correction method to solve these systems of equations for highly ill-conditioned problems with many design variables. Preliminary numerical tests on the potential small disturbance shape optimization problem are promising.

Our plan is to further investigate the above method for applications that are governed by non-linear equations. This approach can be naturally embedded in a SQP formulation of the problem.

This research was conducted in collaboration with A. Battermann and E. Sachs (Universität Trier, Germany).

Large Scale Aerodynamic Shape Optimization

Eyal Arian

The purpose of this work is to develop and apply algorithms which do not require more than a few full solutions of the flow equations to obtain the optimum.

Our approach is to apply approximations in the PDE level to the numerical solution of a practical large scale optimization problem. We are working on shape optimization of a 3D geometry using TLNS3D.

This research was conducted in collaboration with V. Vatsa (NASA Langley).

Electromagnetic Interrogation of Structures

H.T. Banks

The detection and characterization of subsurface damage (cracks, internal corrosions, etc.) is an important problem in aging structures such as airfoils, etc. In collaboration with scientists in the Nondestructive Evaluation Branch at NASA, we have developed computational techniques for inverse problems involving electromagnetic interrogation of structures using superconducting quantum interference devices (SQUIDs).

Our approach is to develop reduced-order model computational methods for Maxwell's equations in a dielectric medium to be used in inverse algorithms. To date, we have developed and tested algorithms based on eddy current interrogation of structures. We use a full Maxwell solver in ANSOFT, which computes time-dependent fields in terms of a vector magnetic potential A in phaser form. Our reduced-order methods are based on Karhunen-Loeve or Proper Orthogonal Decomposition (POD) methods in which we "snapshot" on geometry (length, thickness) of damages in the form of rectangular voices or nontrivial measure "cracks" in 2D.

We have made significant progress on this problem and have tested our ideas with simulated SQUID data for model verification and assessment of the ability to identify and characterize simple damage geometries in a structure. Our findings support the efficacy of this approach and our efforts on more complicated damage paradigms are continuing.

Textbook Multigrid Efficiency in CFD

Achi Brandt

Current multigrid solvers for large-scale steady-state flow problems are still several orders of magnitude slower than the textbook efficiency. The objective of the research is to diagnose the algorithmic obstacles responsible for this situation, separate them through the study of carefully chosen models, and find convenient-to-implement cures for each type of obstacle.

The current research has concentrated on issues related to boundary treatment and flow stagnation points, resulting in the formulation of several new general techniques, such as: a general test to detect boundary troubles; general ways to relax near boundaries and to combine the boundary relaxation with interior line or plane relaxation schemes; modification to the fine-to-coarse interpolation procedure near boundaries whose radius of curvature is comparable to the meshsize (at leading and trailing edges, in particular); and necessary modifications to the relaxation schemes near stagnation points, where the advection coefficients can no longer be regarded as non-principal. Tests with various (mainly incompressible) cases have shown the need for, and the efficiency of, the new techniques.

Further research on increasingly more complicated models is planned, incorporating growing numbers of advance multigrid devices. This research was conducted in collaboration with J.L. Thomas, T.W. Roberts and R.C. Swanson (NASA Langley) and B. Diskin (ICASE).

Distributive Relaxation Technique for Navier-Stokes Equations

Boris Diskin

The current Reynolds-averaged Navier-Stokes solvers with multigrid require on the order of 1500 residual evaluations to converge the lift and drag to one percent of their final values in relatively simple geometries. Complex geometry and/or complex physics simulations need many more residual evaluations to converge, if indeed convergence can even be attained. On the other hand, solutions to elliptic problems attained by the full multigrid process require far fewer, on the 2–4 orders, residual evaluations. Such methods are called demonstrating the textbook multigrid efficiency (TME), which means that the solutions to the governing system of equations are attained in computational work which is a small (say 10) multiple of the operation count in the discretized system of equations. The distributive relaxation technique is one of the most promising techniques for achieving TME in solving complicated computational fluid dynamics (CFD) problems. In framework of this approach, the system is decomposed into separable, many times scalar, pieces that can be treated with optimal methods. In this period, we systematically studied the distributive relaxation technique in application to different settings to incompressible and compressible Navier-Stokes equations.

The two-dimensional Navier-Stokes equations written in nonconservative form were discretized on stretched staggered grids. The compressible flow multigrid algorithm separates two convection components from a system of the two coupled pressure and internal energy equations. Thus, the relaxation process requires a local block 2x2 matrix solution at every grid point. TME was demonstrated for viscous compressible and slightly incompressible flows over a flat plate and for incompressible and compressible wake flows at low (up to 0.7) Mach numbers. All the applications were made for flows at laminar Reynolds number of 10,000. As a step towards development of a general approach, we have repeated the results due to Brandt and Yavneh demonstrating TME for entering incompressible Euler flows at large angles of attack. The authors showed that the standard multigrid V-cycle is not fast enough and suggested using the W-cycle. We derive a series of test problems aimed to separate possible sources of the V-cycle slowdown. Among other findings, we have realized the importance of local boundary relaxations and have demonstrated that the F-cycle is at least as efficient as the W-cycle, but substantially less expensive.

We are going to continue this research on optimal methods for CFD problems. The nearest goals ahead are to attack stagnation flows and compressible flows for all subsonic values of Mach number and angles of attack.

This research was conducted in collaboration with J.L. Thomas (NASA Langley).

Graphical User Interface for a MATLAB System Identification Toolbox

Kerry Huston

Jer-Nan Juang of NASA Langley created his system identification toolbox to allow users to perform system identification on various types of input and output data. The toolbox can perform system ID on time data, pulse response data, or a frequency response function. The programs contained in the toolbox are powerful, but not user-friendly. The purpose of this work was to develop a graphical user interface for the MATLAB toolbox to make it more user-friendly.

The graphical user interface was programmed in MATLAB using the Guide program. The GUI allows the user to enter files, perform system ID, and display results for various types of data. Both MAT and ASCII data file formats are supported by the GUI. The completion of this GUI allows a user unfamiliar with the inner workings of the toolbox to perform system identification.

In the future, system ID toolbox will be utilized to gain a greater understanding of system identification.

Approximation Management in Engineering Optimization

R. Michael Lewis

Great progress has been made in the ability to accurately and faithfully simulate the behavior of physical and engineering systems. However, the enormous computational cost of such simulations makes it impractical to rely exclusively on high-fidelity simulations for the purpose of design optimization. Our objective is to make as much use as possible of models of lower physical fidelity, but lower computational cost, with only occasional recourse to expensive, high-fidelity simulations. This is in keeping with one tenet of nonlinear programming: that one should try to avoid doing too much work when far from an optimizer, since knowledge of the general trends in the objective and constraints can suffice to make progress in the optimization. Moreover, the use of approximation models is in keeping with engineering practice, where models of lower fidelity are widely used in preliminary design to explore the design space.

We have previously developed a trust region approximation management framework (AMF) for the use of general non-quadratic approximations in optimization that insures robust global behavior. The significance of this work is that it is one of the few systematic approaches to the use of non-quadratic approximations and surrogates in nonlinear programming and the first attempt to provide an analytical justification for such strategies.

We have recently completed preliminary numerical tests of the AMF for two aerodynamic optimization problems. The first is a constrained aerodynamic optimization problem for a 3D wing, and the second is a constrained aerodynamic optimization for a 2D airfoil. For each problem, two fidelities of models were implemented based on two levels of mesh discretization.

We have implemented three AMF, one based on an augmented Lagrangian algorithm, one based on an SQP algorithm with an ℓ^1-penalty function as the merit function, and one based on a multi-level nonlinear programming method due to Alexandrov. All of these approaches yielded promising results, with speed-ups of about 3 for the wing problem, and speed-ups of about 2 for the airfoil problem.

We believe we can improve these results by relaxing the degree of optimization performed in the AMF subproblem, which involves the minimization based on the lower-fidelity approximation. We are currently solving this problem to a high degree of precision; however, we need only solve it to a point that yields a suitable improvement in the merit function for the high-fidelity problem. We are currently investigating this question, as well as preparing for tests on more challenging aerodynamic optimization and multidisciplinary problems.

This research was conducted in collaboration with Natalia Alexandrov, Larry Green, Clyde Gumbert, and Perry Newman (NASA Langley).

Probabilistic Methods in MDO

R. Michael Lewis

Many practical design optimization problems include an element of uncertainty. For instance, operating conditions might not be known in advance or may be inherently variable. Material properties may be stochastic. There may also be modeling errors. Moreover, design optimization that ignores uncertainty and probabilistic elements will tend to produce designs that are not robust.

Effective techniques for probabilistic design have been developed in the field of structural reliability. Probabilistic design, in which as one posits that the uncertain quantities are stochastic variables with known (or estimable) statistic, is aspect of the general question of design under uncertainty. We have begun an investigation of whether reliability methods can be applied to problems that are more general than structural analysis.

We have investigated alternative formulations of the most probable point (MPP) calculation in reliability index based design, with the goal of reducing the cost of reliability-based optimization. We have also conducted a preliminary study of applying ideas from reliability-based optimization to multipoint aerodynamic design, where the Mach number and angle of attack are treated as random variables. We began development of model MDO problems with uncertainty. Our first model problem is a coupled aerodynamic, structural, and control problem for supersonic flow in a duct.

This research was conducted in collaboration with Sankaran Mahadevan (Vanderbilt University) and Rakesh Kapania (Virginia Polytechnic Institute and State University).

Pattern Search Methods for Nonlinear Optimization

R. Michael Lewis

Pattern search methods for nonlinear optimization have a number of features that make them attractive for use in engineering optimization. These methods are easy to understand and implement, are scalably parallel, and neither require nor estimate derivatives.

We have previously developed the pattern search algorithms for bound and linearly constraint minimization, and a pattern search method for general nonlinearly constrained optimization guaranteed to possess first-order stationary point convergence. We have more recently undertaken implementation of the new, general classes of pattern search algorithms. This new implementation will allow us to investigate various algorithmic approaches, as well as opportunities for improved computational parallelism.

As part of this work we have studied the asymptotic behavior of pattern search methods and developed an analytical justification for the use of a small pattern size as a practical, yet robust, stopping criterion. Based on comparative tests (conducted at The College of William & Mary) of various pattern search algorithms for unconstrained minimization, we also have evidence that pattern search methods that perform more work per iteration (such as Hooke and Jeeves) are, in fact, more efficient overall. These results are guiding our current development.

This research was conducted in collaboration with Elizabeth Dolan and Virginia Torczon (The College of William & Mary).

Application of Fuzzy Logic to Tune the Parameters of Generalized Predictive Control Algorithm

Chaung Lin

In the generalized predictive control (GPC) algorithm, a control penalty l is introduced to limit the control effort and stabilize the closed loop system. Theory and applications both show that the smaller the value of l, the better the performance of the GPC, which is due to a rather large control magnitude. However, there is a limitation on the control magnitude and so the performance and control magnitude, can be compromised. Another adjustable parameter is the AutoRegressive eXogeneous (ARX) model order. Theoretically, the higher the model order, the better the model output prediction. However, when the model order exceeds a certain number, the model only gains small accuracy improvement and the required computation time becomes rather long, which imposes limitations on the sampling rate using on-line data. It is necessary to adopt a proper model order to optimize computation time. Previously, these two parameters were tuned by trial-and-error, which required considerable effort. Therefore, fuzzy logic rules are developed to tune parameters automatically.

To tune the control penalty, the antecedent part of fuzzy rules has two inputs, performance index (PI) and control index (CI). PI is the ratio of root mean squared of system outputs with control to that without control. CI is the ratio of maximum control magnitude to the allowable control magnitude. The output of fuzzy rules is the adjustment of the control penalty. For tuning the model order, three factors are considered to tune the order. These are the normalized output prediction error (PE), which accounts for model accuracy, change in PE (CPE), and normalized model order (ORD). In order to reduce the fuzzy rule complexity, the fuzzy rules are arranged in hierarchical structure. The first level's input variables are PE and CPE and output is an order multiplier, which is called OUT1. The second level's inputs are ORD and OUT1 and the output is a multiplier that adjusts the result of the first level.

The performance of the developed fuzzy logic method has been demonstrated in simulations. The results are similar to those obtained by trial-and-error. The future plan is to implement the rules on the real-time control problem.

Robust Multigrid Methods for the 3D Incompressible Navier-Stokes Equations

Ignacio M. Llorente

Alternating-plane smoothers have been found to be robust and highly efficient smoothers for the solution of anisotropic elliptic operators even for multi-block grids. We continue the work on alternating-plane smoothers by studying the applicability of these methods as smoothers for the Navier-Stokes equations. Several methods that combine implicit point, line, or plane relaxation with partial and full coarsening have been proposed in multigrid literature to solve anisotropic operators and achieve robustness when the coefficients of a discrete operator can vary throughout the computational domain (due to grid stretching or variable coefficients). These methods have been always studied using the scalar diffusion-convection equation. In this way, conclusions could be obtained to choose the best method for more complicated mathematical models (real models as the incompressible Navier-Stokes equations). However, they have been rarely applied to solve such more complicated models. Our main goal is to compare those robust alternatives for the incompressible Navier-Stokes equations in a single-block grid with stretching.

A FAS multigrid code to solve the incompressible Navier-Stokes equations is under development. The discrete equations are obtained using a finite volume approach in a staggered grid. For the convective terms, a first-order Upwind method is used. Second-order accuracy is achieved with a defect-correction procedure based on a QUICK scheme. The staggered arrangement was implemented due to its accuracy, stability, and conservation properties. The smoothing operator is a cell-implicit Symmetric Coupled Gauss Seidel method, where the momentum and continuity equations are relaxed in a coupled manner. This smoother has been reported to be robust and stable because of the coupling between the equations. The code allows any combination of cell or block relaxation with partial or full coarsening. The 2D problems in each plane can also be solved by any combination of cell or line relaxation with full coarsening and the 1D problems are again solved by multigrid. The work performed will demonstrate the robustness and efficiency of alternating-plane smoothers even for complicated PDEs.

We intend to continue work on efficient parallel multigrid methods for block-structured applications. In particular, we will compare simultaneous and distributive relaxation for the incompressible Navier-Stokes equations. Following, the study will be extended to compressible Navier-Stokes problems.

This research was conducted in collaboration with R.S. Montero (Complutense University of Madrid) and N.D. Melson (NASA Langley).

A Highly Parallel Multigrid Solver for 3D Upwind-biased Discretizations of Convection-dominated Problems

Ignacio M. Llorente

Standard multigrid algorithms are known to be highly efficient in solving systems of elliptic equations. However, multigrid algorithms fail to achieve optimal grid-independent convergence rates for nonelliptic problems. In many cases, the nonelliptic part of a solver is represented by the advection operator. Downstream marching, when it is viable, is the simplest and most efficient way to solve this operator. However, in a parallel setting, the sequential nature of the marching degrades the efficiency of the algorithm. The aim of this research is to evaluate and analyze an alternative highly parallel multigrid solver for a 3D advection operator using semicoarsening and a colored plane-implicit smoother.

We are currently developing a multigrid code to solve the nonlinear 3D diffusion-convection equation on cell-centered grids in convection-dominated regimes. The multigrid method is based on four-color plane-implicit smoothers combined with semicoarsening. Well-balanced explicit correction terms are added to coarse grid operators to maintain on coarse grids the same cross-characteristic interaction as on the target (fine) grid. This approach allows an efficient solution (grid independent convergence rate and highly parallel implementation) for convection-dominated problems even in cases of nonalignment, i.e., when the grid lines do not align with the characteristic direction.

In order to study the parallel efficiency and scalability of the approach, the code will be implemented on the Cray T3E, SGI Origin 2000, and the Coral system using MPI.

This research was conducted in collaboration with B. Diskin (ICASE).

Evaluation and Analysis of the Parallelization of a Multigrid 3D Incompressible Navier-Stokes Solver on the Coral System

Ignacio M. Llorente

Computer clusters have emerged as a cost-effective solution for parallel scientific computation. However, the achievable performance of Beowulf systems, built using personal computers interconnected by switched fast ethernet, is not clear when running intensive numerical applications. Their lower cache memory size (16KB L1 and 512KB L2), network bandwidth (100Mbit/s), and floating-point performance may reduce the achievable percentage of their theoretical peak performance. Related to the communication performance, we have previously reported that the achievable bandwidth obtained in a real application depends not only on the amount of communication but also on how the data are structured in the local memories, due to their impact on the send and receive overheads, and how the network topology is exploited due to contention problems. Bandwidth reduction, due to non-unit-stride memory access, and contention in the network may affect the optimal decomposition of structured CFD applications. The aim of this research is to evaluate the effect of the local memory use and the communication network exploitation on message sending in the Coral system and study how these results affect the optimal decomposition of structured applications.

We have performed communication tests without network contention to establish the achievable bandwidths, and then, we have modified this base experiment by increasing the contention of the network and by decreasing the spatial locality properties of the messages. Preliminary results show similar bandwidth to the SGI Origin 2000 when the messages are non-contiguous in memory (around 8 Mbytes/s). We are currently analyzing these results, taking into account the underlying architecture (memory hierarchy, microprocessor, and interconnection network). A parallel multigrid solver of the 3D incompressible Navier-Stokes equations in a staggered grid is under development. The MPI code will allow partitioning in any direction in order to study the optimal decomposition of structured applications on the Coral system. The code implements a FAS algorithm with a cell-implicit Symmetric Coupled Gauss Seidel smoother, where the momentum and continuity equations are relaxed in a coupled manner. This study will provide guidance for efficient programming in the Coral system.

The performance results will be compared with those obtained on some current parallel computers (Cray T3E and SGI Origin 2000) considering not only performance issues but also the performance-cost ratio.

This research was conducted in collaboration with M. Prieto-Matias (Complutense University of Madrid).

Spatial Structure of Optimal Flow Control

Josip Lončarić

Designing distributed control systems begins with the sensor/actuator placement problem. While in some situations discrete search of combinatorial complexity seems unavoidable, continuum problems suggest solving a related question. *If one could sense everything and actuate everywhere, what should one do?* The answer to this question has polynomial complexity (of order N^3 where N is the number of state variables) and can serve as the initial effectiveness filter capable of rejecting a large portion of the design search space. This favorable situation can have several causes depending on the base flow pattern. Our aim is to develop efficient numerical procedures to solve this problem for flows in moderate Reynolds number regimes.

In an earlier work, we developed a rational approximation of the optimal feedback kernel for unsteady Stokes flow. For the flow around a cylinder, this approximation was proven to perform within 0.026 percent of the exact optimum even in the worst case. Using the vorticity representation in conformally mapped geometries, this approximation is decomposed into the analytic free space solution and a boundary term which can be evaluated numerically. This procedure was applied to the NACA 0015 wing. The results demonstrate a significant contribution of the boundary to the control effort. We investigated the Stokes preconditioner for the nonzero base flow case. As a first step, we tested this approach on a shear flow where the Fourier transform in the streamwise direction can be used to simplify the problem. We obtained a number of reduced order models and the corresponding optimal control operators accounted for about 90 percent of the full numerical solution.

The insight gained in this study will provide guidance for the development of numerical schemes for the full NACA 0015 wing case at moderate Reynolds number flows.

The Coral Project

Josip Lončarić

The cost of developing complex computer components such as CPUs has become so high that scientific applications alone cannot carry the full burden. In the future, scientific computing will have to use mass market leverage to overcome the cost barrier. A cost-effective alternative to high-end supercomputing was pioneered by Beowulf, a cluster of commodity PCs. By now, high performance Beowulf clusters can be built using fast commodity PCs and switched Fast Ethernet. We want to explore the benefits and the limitations of this approach, based on applications of interest to ICASE.

The initial phase of the Coral project, consisting of 32 Pentium II 400 MHz nodes and a dual-CPU server, demonstrated aggregate peak performance in excess of 10 Gflop/s, with sustained performance on CFD applications of about 1.5 Gflop/s. In order to provide a richer environment for further experimentation, a dual-CPU configuration was chosen for the second phase of the Coral project. We have added 32 Pentium III 500 MHz processors (configured as 16 dual-CPU nodes) and two dual-CPU file servers. The resulting system contains 64 compute CPUs with an aggregate of 20 GB of RAM and 440 GB of local disk space. The three dual-CPU servers provide an additional 1.5 GB of RAM and 240 GB of disk space.

We had to resolve a number of performance problems with the new dual-CPU nodes, and tuning continues. Nonetheless, Coral has an excellent price/performance ratio, almost an order of magnitude better than an equivalent proprietary supercomputer design. This conclusion is based on our experience with a variety of applications, ranging from coarse-grained domain decomposition codes to communication-intensive parallel renderers.

For a computer system with this much memory and disk capacity, the ability to move data around within the cluster, as well as to and from other systems, becomes an important consideration. In conjunction with the second phase of the Coral project, we have designed a small gigabit ethernet testbed which will connect the two large file servers, two of the dual-CPU compute nodes, and one of ICASE's graphics workstations. This configuration will be used to evaluate price/performance tradeoffs of various gigabit ethernet components, and to determine the utility of gigabit ethernet for file transfers, cluster I/O, and interactive visualization.

We plan to finish refining the performance of the new nodes and to update the old nodes to Linux kernel 2.2.12. The cluster will be used to develop and run research codes of interest to ICASE and NASA Langley, and to evaluate price/performance tradeoffs among various hardware, software, and networking configurations.

This research was conducted in collaboration with T.W. Crockett, P. Mehrotra, S. Zhou, and M.D. Salas (ICASE).

Active Shielding and Control of the Environmental Noise

Josip Lončarić

Rejection of exterior noise caused by periodic sources such as propellers or turbines would significantly enhance passenger comfort and reduce noise fatigue on long flights. Passive sound absorbing materials help at high frequencies, but to be effective below about 1 kHz their weight penalty becomes significant. Active noise control can reduce low frequency noise with less weight penalty. We present the mathematical foundations of a new active technique for control of the time-harmonic acoustic disturbances.

Unlike many existing methodologies, the new approach provides for the exact volumetric cancellation of the unwanted noise in a given predetermined region of space while leaving those components of the total sound field deemed as friendly unaltered in the same region. Besides, the analysis allows us to conclude that to eliminate the unwanted component of the acoustic field in a given area, one needs to know relatively little; in particular, neither the locations nor the structure of the external noise sources need to be known. We constructed the general solution for the aforementioned noise control problem. The apparatus used for deriving the general solution is closely connected to the concepts of generalized potentials and boundary projections of Calderon's type. To prove that the new technique is appropriate, we thoroughly worked out a simple two-dimensional model example that allows full analytical consideration, including optimalization of the control effort.

In order to develop numerically computable solutions, we plan to describe the discrete framework for the noise control problem parallel to the continuous one. This discrete framework is obtained using the difference potentials method; in the future it is going to be used for analyzing complex configurations that originate from practical designs. Once we have computed the solution for a particular configuration, we intend to investigate the possibilities of optimizing it according to the different criteria that would fit different practical requirements. We expect to discuss the applicability of the technique to quasi-stationary problems, future extensions to the cases of the broad-band spectra of disturbances, as well as other possible applications which may include different physics, such as electrodynamics, and different formulations of the boundary-value problems, such as scattering.

This research was conducted in collaboration with S.V. Tsynkov (Tel Aviv University, Israel).

Analysis of Lattice Boltzmann Equation: Dispersion, Dissipation, Isotropy, Galilean Invariance, and Stability

Li-Shi Luo

The method of the lattice Boltzmann equation (LBE) has been applied to various areas in computational fluid dynamics. The dynamics of the lattice Boltzmann equation evolves on a highly symmetric lattice space, which is usually square in 2D and cubic in 3D. Therefore, the transport properties of the LBE method are dictated by the symmetry of the underlying lattice. The present work proposes a general procedure to study the dispersion, dissipation, isotropy, Galilean invariance, and stability of the LBE method.

We first propose a generalized LBE nine-velocity model in two-dimensional space. The model is constructed in moment space instead of velocity space, and possesses a maximum number of adjustable parameters allowed by the freedom of the LBE model. We then analyze the dispersion equation of the linearized evolution operator to obtain the generalized hydrodynamics of the model (the wave vector $\mathbf{k}$-dependence of the transport coefficients). We study the dispersion, dissipation, isotropy, Galilean invariance, and stability of the LBE model, and optimize the properties of the model by tuning the adjustable parameters in the model according to our linear analysis. We find that the proposed model is superior to the popular BGK

LBE model in terms of stability, isotropy, and Galilean invariance. Various LBE models are also analyzed and compared with each other.

A paper entitled "Theory of the lattice Boltzmann method: Dispersion, dissipation, isotropy, Galilean invariance, and stability," authored by Pierre Lallemand and Li-Shi Luo, is under preparation.

The present work is a result of collaboration with Pierre Lallemand (Laboratoire ASCI-CNRS, Université Paris-Sud (Paris XI Orsay), France). The present work has been funded by NASA Langley Research Center under the program of "Innovative Algorithms for Aerospace Engineering Analysis and Optimization."

We intend to extent our work to 3D LBE models or other more complicated LBE models.

Probabilistic Methods for Multidisciplinary Optimization Under Uncertainty
Sankaran Mahadevan

The planning and design of engineering systems is always done under conditions of uncertainty regarding operating conditions, system behavior, etc. Traditionally, empirical safety factors have been used in designing engineering systems for high reliability. The estimation of reliability has been mostly done through testing, which is feasible for small devices and equipment, but not for large systems. Recent efforts at developing physics-based probabilistic computational methods to predict and design for reliability have been successful in the structures discipline. In this research, we are investigating the extension of these methods to aerodynamic systems and developing new formulations that allow reliability analysis and uncertainty-based design of coupled, multi-disciplinary systems.

In probabilistic reliability analysis, the uncertainties are quantified through stochastic (random) variables with appropriate statistical information. Reliability is computed as the integral of the joint probability density function over the domain of satisfactory performance. Multi-dimensional integration and Monte Carlo simulation are two strategies for this computation that are very expensive. Therefore, in this research, we pursue analytical first-order and second-order approximation methods that convert the problem into one of optimization, of finding the minimum distance of the limit state from the origin in an equivalent space of standard normal variables. For single discipline structures problems, a quasi-Newton recursive formula has proved successful. For coupled multi-disciplinary problems where analysis codes in different disciplines usually exist separately, a new formulation is necessary. We have developed several alternative formulations that add or subtract auxiliary optimization variables (or degrees of freedom) in a manner that is applicable to multi-disciplinary problems. We have developed computer codes incorporating these methods and implemented them for small demonstration problems. Based on this, we have also developed a new formulation for uncertainty-based design optimization of multi-disciplinary systems.

We plan to apply these methods to large multi-disciplinary problems that involve aerodynamics, structures, heat transfer, and controls disciplines. This effort will include the characterization of different sources of uncertainty and their incorporation in system analysis and design. In addition, we plan to investigate the convergence characteristics of the bi-level problem of design optimization under probabilistic constraints. This work is done in close collaboration with R.M. Lewis (ICASE) and R. Kapania (Virginia Polytechnic Institute and State University).

Large Eddy Simulation Using a Parallel Multigrid Solver
Dimitri J. Mavriplis

The failure to develop a universally valid turbulence model coupled with recent advances in computational technology have generated a greater interest in the large-eddy simulation approach for computing flows with large amounts of separation. This approach involves resolving the large-scale unsteady turbulent eddies down to a universally valid range in the hope of yielding a more generally valid simulation tool. The purpose of this work is to develop a large-eddy simulation capability based on an existing unstructured grid Navier-Stokes solver. The use of unstructured grids, which facilitates the discretization of complex geometries and adaptive meshing techniques, is expected to enhance the flexibility of the resulting simulation capability.

The first step to developing a large-eddy simulation capability involves the extension of the currently existing parallel unstructured multigrid steady-state Reynolds-averaged Navier-Stokes flow solver to an unsteady Reynolds-averaged Navier-Stokes flow solver. This is achieved by discretizing the time derivative using a three-point backwards stencil in time, and solving the nonlinear equations at each time-step with the steady-state unstructured multigrid solution algorithm.

The extension of the steady-state solver to handle unsteady flows has been implemented, and validation is currently underway. The initial validation case consists of the unsteady flow over a circular cylinder, with emphasis on obtaining the correct Strouhal number. A grid convergence study (in time and space) will be performed to determine the effect of resolution on overall accuracy. These computations are being performed on the ICASE PC cluster, Coral.

Once a validated unsteady Reynolds-averaged Navier-Stokes solver is in hand, the implementation of a subgrid scale model will be pursued. This model will be validated by computing several standard decaying turbulence test cases, and will then be applied to cases of aerodynamic interest.

This research was conducted in collaboration with Juan Pelaez (Old Dominion University).

Scalability of an Unstructured Multigrid Flow Solver Using Various Programming Models

Dimitri J. Mavriplis

Under the ASCI program, the need to simulate very large problems on massively parallel machines has focused attention on the relative merits of various parallelization strategies. The implementations of interest include the message-passing interface (MPI) library and the OpenMP parallelization standard. While OpenMP has been devised for use on shared-memory architectures, MPI can be implemented on any type of architecture, but is primarily attractive for distributed memory architectures. Furthermore, the emergence of clusters of shared memory multiprocessors as a viable scientific computing architecture has resulted in interest in hybrid OpenMP-MPI programming models.

The goal of this work is to develop a parallel unstructured mesh flow solver using both MPI and OpenMP, and to examine the scalability of the solver on various architectures using the different programming models either alone, or in the combined hybrid mode.

An existing unstructured multigrid flow solver which was originally parallelized using the MPI approach has been upgraded to include an OpenMP parallelization strategy. OpenMP parallelization is achieved within the existing MPI parallel structure, in a two-level nested fashion. In this manner, the same code can be run in pure MPI mode, in pure OpenMP mode, or in the combined OpenMP-MPI mode, which would seem appropriate for clusters of shared memory multiprocessors.

Preliminary results indicate that on "physically" shared memory machines, such as a Cray SV1, the OpenMP and MPI approaches yield similar scalability. On shared memory machines with "physically" distributed memory, such as the SGI Origin 2000, the MPI approach appears to scale better at high processor counts.

To date, only very preliminary results have been obtained. More benchmarking on various architectures is to be performed, and a closer look at the memory placement issues is required for machines such as the SGI Origin 2000. Finally, the scalability of a hybrid two-level OpenMP-MPI application running on a large cluster of shared memory multiprocessors is to be investigated and contrasted with that of a purely MPI-based approach.

Alternative Techniques for Implicit Residual Smoothing

Cord-Christian Rossow

Implicit residual smoothing is a means to extend the local stability range of a basic time-stepping scheme. Most frequently implicit residual smoothing is used in combination with explicit Runge-Kutta time-stepping schemes, and together with multigrid acceleration, fast and efficient codes have been developed. However, implicit residual smoothing has mainly been restricted to structured meshes. This is due to the fact that the cell aspect ratio is explicitly needed to derive the smoothing coefficients for the different coordinate directions. The present work is directed towards the development of an implicit smoothing technique which does not directly require this information. Based on the MAPS (Mach number based Advection Pressure Splitting) spatial discretization, a smoothing technique is to be derived which, similarly to the original aspect-ratio based scheme, employs only the solution of scalar equations, however the computation of cell aspect ratios shall be avoided.

Several efforts during the last years have been focused on the development ofdiscretization methods that combine the accuracy of flux-difference splittings with the robustness of flux vector splittings. One recent contribution to this class of hybrid flux splittings is the MAPS scheme. Further research revealed that the scheme is very similar to the Roe flux-difference splitting, with the exception that no intermediate state needs to be computed, and no entropy condition is required. It was found that in the original MAPS formulation only the compressible terms of the Roe-scheme are retained. Including the incompressible terms of the Roe-scheme into the MAPS formulation extended MAPS to incompressible flows, yielding then the MAPS+ discretization.

In the present research, the connection of the MAPS/MAPS+ discretizations with the Roe-scheme was further exploited. The first topic investigated was the combination of the MAPS/MAPS+ spatial discretizations with an implicit formulation for the time integration. For implicit schemes the flux Jacobians need to be evaluated, which is well-established for the Roe-scheme. Due to the similarity of MAPS/MAPS+ and Roe discretizations, in a first step the discretization of the left-hand side was based on the Roe-scheme. In order to pave the way for an implicit smoothing technique, the implicit formulation was incorporated into the framework of an explicit Runge-Kutta scheme. This was done similarly as j-line preconditioning is usually implemented into Runge-Kutta schemes. The resulting implicit, linear system for the residuals is solved at each stage of the Runge-Kutta scheme via simple Jacobi iteration. For the solution of the Navier-Stokes equations,it was furthermore found necessary to include an ADI-scheme as a preconditioner for the Jacobi iteration. It was essential to perform the line-implicit solution across the wake region of the airfoil. The Jacobi iteration was then used to reduce the factorization error. In combination with multigrid acceleration, with implicit scheme convergence rates of 0.7 for inviscid flows, rates of 0.82 for turbulent viscous flows around airfoils were obtained.

The implicit formulation yields favorable convergence rates, however, a large amount of storage is required for the flux Jacobians. Additionally, the solution of the implicit system requires a considerable amount of work. Therefore, in the second phase of the research, simplifications of the classical implicit formulation were sought to derive an implicit residual smoothing technique, which is as simple as common implicit residual smoothing schemes, but also generally applicable. First, the original MAPS formulation without incompressible terms was regarded. Treating the pressure terms in the momentum equations explicitly, scalar implicit equations for the conservative residuals could directly be derived from the MAPS discretization. These implicit, scalar equations were incorporated into the Runge-Kutta scheme in the same way as described before for the complete implicit system. With this implicit smoothing technique, CFL numbers comparable to standard implicit smoothing schemes were obtained. The technique derived from the MAPS scheme does not require any cell aspect ratio information, since it was directly derived from the spatial discretization. Further testing of the MAPS smoothing technique however revealed, that the smoothing could only be combined with the original MAPS scheme. Adding the incompressible terms of the MAPS+ scheme resulted in divergence. Furthermore, low Mach number preconditioning could not be applied with this smoothing technique. To resolve these problems, the pressure terms arising in the incompressible terms needed to be included. These terms however lead to a strong coupling of the equations, thus the desirable property of solving only scalar equations in the implicit smoothing step would be lost. It was found that the influence of the incompressible terms could approximately be taken into account by introducing a term comprised of the speed of sound multiplied by the scaling factor of the incompressible terms. With this modification, the MAPS+ discretization could also be used in combination with the new implicit smoothing. Furthermore, low Mach number preconditioning could be applied by using the preconditioned speed of sound in the implicit smoothing. Note that this additional term is added to each of the scalar equations, thus no coupling of the equations is generated, and the simple structure of the smoothing scheme could be retained. It was also found that the explicit pressure term could now be dropped. With this simplification, the numerical effort for the new smoothing technique is the same as for common residual smoothing techniques. The numerical efficiency observed on structured meshes therefore, was also similar to usual techniques. The new smoothing technique were then implemented into an unstructured code. Preliminary computation of inviscid and viscous flow on hybrid meshes showed the possible increase in CFL number and convergence.

The implementation of the newly developed smoothing technique yielded an increase in CFL number as observed with usual smoothing techniques. When implemented into the unstructured code, the increase in CFL number was also observed. However, in combination with multigrid acceleration, after a fast decrease of the mean residual by about four orders of magnitude, the asymptotic convergence degraded in contrast to the structured code. Thus, the multigrid implementation of the unstructured code needs to be checked. Furthermore, it was found that the reconstruction procedure used in the unstructured code needs refinement to allow fast convergence to steady state. These are issues which will be dealt with independently from the research described above.

High-order Discontinuous Galerkin Method and WENO Schemes

Chi-Wang Shu

Our motivation is to have high-order non-oscillatory methods for structured and unstructured meshes, which are easy to implement on parallel machines. The objective is to develop and apply high-order discontinuous Galerkin finite element methods and weighted ENO schemes for convection dominated problems. The applications will be problems in aeroacoustics and other time-dependent problems with complicated solution structure.

Jointly with Harold Atkins at NASA Langley, we are continuing the investigation of developing the discontinuous Galerkin method to solve convection-dominated convection diffusion equations. Emphasis for this period is put upon studying the eigenvalue structures of the discontinuous Galerkin operator on the diffusion terms. The eventual objective is to design efficient implicit time discretization techniques for such operators. Jointly with Jian-Guo Liu of University of Maryland, we have developed a Discontinuous Galerkin method for 2D incompressible Euler and Navier-Stokes equations in the vorticity formulation. Numerical experiments indicate good resolution of the method near singular solutions. Concerning weighted ENO schemes, jointly with Changqing Hu and Jing Shi of Brown University, we are studying a new technique to handle negative linear weights. Preliminary numerical results are very promising.

Research will be continued on high-order discontinuous Galerkin methods and weighted ENO methods and their applications.

Upwind High-resolution Factorizable (UHF) Schemes for the Equations of Fluid Flow

David Sidilkover

In order to develop an optimally efficient solver for the general equations of fluid flow a new discretization scheme is needed. It has to imitate the mixed type of the PDEs, i.e., to be *factorizable* — to make the different factors of the equations distinguishable on the discrete level. In addition to that it has to possess the shock-capturing property, i.e., to be conservative and high-resolution. The scheme also has to be h-elliptic, i.e., its residuals should be sensitive to the high-frequency error components. There was no discretization that satisfies all these requirements. The main objective of this research is to develop one.

An upwind high-resolution factorizable (UHF) scheme for the Euler equations in two dimensions has been constructed. The factorizability property can be utilized for the purpose of designing an optimal solver through a "projection-type" relaxation. This relaxation procedure relies on the auxiliary potential and stream-function variables. Another important implication of the factorizability property is that the scheme should not lose accuracy for the low Mach number flow. In other words, the proposed approach essentially provides the long-awaited unification of the compressible and incompressible flow solvers. The new method is also expected to produce results of higher fidelity than the standard methods, especially for the vortical flows.

The current work is devoted to extending the scheme/solver to the three-dimensional case and general body-fitted grids. Also, one of the important future directions is to address the transonic/supersonic cases.

Towards the Practical Application of the UHF Scheme

David Sidilkover

The development of the UHF scheme for the case of Cartesian grids has been completed. The next task would be to apply these ideas to the "real world" problems. The first step in this direction is to address the issue of complex geometries.

This is done through extending the UHF scheme for the generalized coordinates. The vector formulation of the scheme appears to be of crucial importance to achieve this while making sure that the factorizability property is preserved. The control volume approach is used. This makes it very easy to implement the UHF scheme within the existing aerodynamic codes. The initial numerical experiments clearly demonstrate the robustness and accuracy of the new methodology. Uniform convergence rates have been achieved for the subsonic flow. One of the unique features of the new scheme is that there is neither loss of efficiency nor loss of accuracy in the incompressible limit.

The future work should address 3D and transonic/supersonic cases.

This research was conducted in collaboration with T.W. Roberts and J.L. Thomas.

On the Application of Global Boundary Conditions with the New Generation of Flow Solvers

David Sidilkover

Previous work indicates that the global highly accurate artificial boundary conditions implemented along with a multigrid flow solver may lead to major improvements in the overall algorithmic efficiency. These improvements are due to the substantial reduction of the domain size with no accuracy loss, better robustness, and faster multigrid convergence. However, the flow solvers considered so far were the standard ones which are far from being optimally efficient. The objective of this research is to study the application of the global boundary conditions with the new generation of the flow solvers, which are optimal.

The new generation of flow solvers is based on factorizing the flow equations, i.e., distinguishing between the factors of different type. This simplifies the derivation of the global boundary conditions, since each factor can be treated separately for this purpose. The boundary condition accurately takes into account the structure of the far-field solution; for simple settings this is done directly using Fourier transform and separation of variables, for more general settings the difference potentials method (DPM) is used. The issue of relaxing properly the discrete equations together with the new global boundary conditions requires special attention in order to maintain the optimal convergence rates of the overall solver. This was the main subject of the current research. As an outcome of the project, we expect the multifold reduction of the size of the computational domain, while preserving the accuracy of the solution and the optimal convergence rates of the solver.

Currently, the algorithm based on the separation of the elliptic equation for the pressure is under study. In the future we plan on applying the global boundary conditions in conjunction with the new upwind high-resolution factorizable (UHF) scheme.

This research was conducted in collaboration with S.V. Tsynkov and T.W. Roberts.

Global Artificial Boundary Conditions for Aerodynamic and Aeroacoustic Computations

Semyon Tsynkov

Many problems in scientific computing, including those that present immediate practical interest, e.g., problems of acoustic radiation/propagation/scattering, are formulated on infinite domains. Therefore, any numerical methodology for solving such problems has to be supplemented (or, rather, preceded) by some technique that would lead to a finite discretization. Typically, the original domain is truncated prior to the actual discretization and numerical solution. Subsequently, one can construct a finite discretization on the new bounded computational domain using one of the standard techniques: finite differences, finite elements, or other. However, both the continuous problem on the truncated domain and its discrete counterpart will be subdefinite unless supplemented by the appropriate closing procedure at the external computational boundary. This is done by using artificial boundary conditions (ABCs); the word "artificial" emphasizing here that these boundary conditions are necessitated by numerics and do not come from the original physical formulation.

At the current stage of the aforementioned project, we are focusing on the development of highly accurate unsteady ABCs for the numerical simulation of waves propagating with finite speed over infinite domains. The major difficulty is the nonlocal character of these boundary conditions in both space and time. Typically, the spatial nonlocality is tolerable, while the temporal one is prohibitively expensive. To overcome this difficulty, we make use of the presence of lacunae in the solutions of some linear hyperbolic PDEs in oddly dimensional spaces. Based on this property of the solutions, we construct a special lacunae-based algorithm for the long-term numerical integration of hyperbolic PDEs, this algorithm does not accumulate error and has fixed non-growing expenses per time step. Incorporating this lacunae-based algorithm into the general procedure of constructing the ABCs based on Calderon's projections and difference potentials, we gain the restriction of the temporal nonlocality of the boundary conditions. Some initial calculations for the three-dimensional wave equation excited by a compactly supported stationary source have already provided the encouraging results. The code for the wave equation excited by a moving source is in the development stage. ICASE Report No. 99-23 is the first publication on the subject, other publications are forthcoming.

Future research in the framework of this project will primarily concentrate the development of unsteady ABCs for the actual problems in acoustics, including the advective case, and electromagnetics. The project is a collaborative effort with V. Ryaben'kii, V. Turchaninov, and H. Atkins. The project is supported by the Director's Discretionary Fund.

Adjoint Error Estimation and Grid Adaptive Criteria for Accurate CFD Predictions of Integral Outputs

David A. Venditti

Accurate predictions of key integral quantities such as lift or drag are often of primary concern in CFD simulations of complex engineering problems. In an effort to improve the accuracy of these quantities and reduce computational costs, a multilevel error estimation and grid adaptive strategy is being developed and investigated. The error estimation procedure is based on a discrete adjoint formulation that relates the simulation error in an integral quantity to the local residual errors of both the primal and adjoint solutions. The magnitudes of the local error contributions can be used as criteria in a grid-adaptive strategy designed to yield specially tuned grids for accurately estimating the error in the chosen functional. The objective of this work is to implement the procedure in FUN2D (A Fully UNstructured 2D flow solver developed by W. Kyle Anderson) and apply it to several inviscid and viscous test problems.

The procedure has already been implemented and tested using a standard second-order finite volume discretization of the quasi-1D Euler equations. A doubling in the accuracy of the integrated pressure along a converging-diverging duct was achieved after applying the adaptive strategy and error estimation procedure to the problem. Preliminary 2D inviscid and viscous results were also obtained with FUN2D. Using the error estimation and adaptive grid method, significant improvements in the accuracy of the computed lift and drag were achieved for several different single-body configurations including transonic flows with shocks.

We intend to implement more aggressive grid-adaptive capabilities including anisotropic grid adaptation suitable for viscous flow applications. Furthermore, we intend to apply the procedure to more complicated geometries including multi-element airfoil configurations and, ultimately, to 3D aircraft configurations.

This research is being conducted in collaboration with David L. Darmofal (Massachusetts Institute of Technology) and W.K. Anderson (NASA Langley).

A Gas-kinetic Scheme for Two-phase Flows

Kun Xu

The study of liquid/gas phase transition and interface movement is very important in science and engineering, which could find wide applications in environment, geophysics, and engineering. Current existing approaches for the two phase flows are mainly based on the front tracking methods, where the liquid/gas interfaces are explicitly constructed. In order to simplify the interface tracking mechanism and develop an efficient interface capturing scheme, we have solved the hyperbolic-elliptic equations for the two phase flows, where the van der Waals equation of state automatically takes care of the phase transition between liquid and gas. The present work is about the development of a gas-kinetic scheme for the mixed-type system.

The gas-kinetic two-phase flow solver is based on the numerical solution for the BGK-type kinetic model. With the appropriate choices of all parameters in the Maxwellian equilibrium distribution function, the hyperbolic-elliptic equations with van der Waals equations of state can be precisely recovered. Once the BGK-type model is obtained, the integral solution of the model is used in the construction of the numerical scheme. Due to the coupling of the spatial and temporal variations of the particle distribution function, the evolution of the two-phase flow around each cell interface is determined, from which the numerical fluxes across the cell interface can be accurately evaluated. The scheme is very robust in comparison with other well-defined shock capturing schemes in the application to the mixed-type system. Because of the intrinsic diffusive and dissipative mechanism in the kinetic scheme, the Maxwell's equal-area construction is automatically accomplished, where the liquid and gas densities settle down to the physically correct values. Since the fluid stays in an unstable state in the elliptic region, any numerically "averaged" or "dissipated" state around the material interface must evaporate to the gas or condensate to the liquid. As a result, the fluid interface can always be kept sharp even after thousands of time steps. This kind of interface capturing mechanism based on the physics, i.e., the van der Waals equation of state, seems more reasonable than any other existing artificial interface sharpening techniques. The current method has been successfully applied to the phase transition and liquid/gas interface movement problems, such as the liquid drop collisions.

We plan to further develop the scheme by including energy equations, such as the latent heat release, and apply the scheme to the boiling heat transfer problem. Due to the simplicity of the current approach, the simulation of hundreds or even thousands of bubbles will become possible.

A High-order Accurate Method on a Moving Grid Adapted to the Solution

Nail Yamaleev

Grid adaptation has now become widespread for solving multidimensional partial differential equations in arbitrary-shaped domains. One of the most important problems associated with the adaptive grid generation is an essential effect of the grid point distribution on error in the numerical solution. Until the present time little attention has been paid to the fact that the concentration of grid points in regions which most influence the accuracy of the numerical solution may at the same time introduce an additional error due to the grid non-uniformity. The main objective of this research work is to construct an optimal coordinate transformation so that the leading term in the truncation error of a finite difference approximation is minimized that provides superconvergent results on the optimal grid.

A new grid adaptation strategy, which minimizes the truncation error of a p-order finite-difference approximation is proposed. The main idea of the method is based on the observation that the local truncation error associated with the discretization on nonuniform meshes can be minimized if the interior grid points are redistributed in an optimal sequence. The method does not explicitly require the truncation errorestimate and at the same time it allows one to increase the design order of approximation by one so that the same finite-difference operator reveals superconvergence properties on the optimal grid. Another very important characteristic of the method is that if the metric coefficients and the firstderivative of the function itself are approximated by the same hybrid scheme, then the optimal grid generator can be employed in the whole computational domain independently of points where the hybrid scheme switches from one approximation to another. Generalization of the present method to multiple dimensions is presented. Numerical calculations of different one-dimensional and two-dimensional finite-difference approximations demonstrate the performance of the method and corroborate the theoretical results.

We plan to apply the present method to calculate both steady and essentially unsteady flows with and without strong discontinuities. This research was conducted in collaboration with M. Carpenter and J. Thomas (NASA Langley).

A Non-dissipative Staggered Fourth-order Accurate Explicit Finite Difference Scheme for the Time-domain Maxwell's Equations

Amir Yefet

We have provided stable finite difference operators to implement boundary/interface conditions in a fourth-order accurate extension of the Yee scheme. Importantly, the convergence rate exhibited by the scheme in the absence of boundaries is preserved in their presence; fourth-order is also obtained in the presence of discontinuous dielectric properties.

For open problems, one would have to employ an appropriate boundary condition to truncate the computational domain which is embedded in free space. The application of a reflectionless sponge layer as an absorbing boundary condition for this scheme has been explored, and we note that our extrapolation and one-sided difference operators may allow for the application of recently developed exact absorbing boundary conditions.

It is our hope that the work herein will provide the incentive to researchers in CEM to revise their opinion of the scheme herein, and to direct efforts towards its further development in order to provide a high-order replacement of the Yee scheme.

We intend to implement this new scheme on usable problems in the future. We also plan to extend this scheme to complex domains.

This research was conducted in collaboration with Peter Petropoulos (New Jersey Institute of Technology).

Evolution of Disturbances in Three-dimensional Boundary Layers

Ponnampalam Balakumar

Our objective is to compute the evolution of disturbances generated from a harmonic point source in general three-dimensional boundary layers.

In general three-dimensional boundary layers, the mean boundary layer profiles vary in all three directions. In earlier work, the stability characteristics of a supersonic boundary layer over a sharp cone at an angle of attack were investigated as a 2D eigenvalue problem. The results showed that due to the variations of the meanflow in the azimuthal direction, the eigenmodes are clustered into isolated, confined regions. It was also observed that there exist several eigenvalues for a fixed axial wavenumber and all of them remain close together. The next step is to investigate the evolution of disturbances which are induced by harmonic point sources located at the wall in a three-dimensional boundary layer. The disturbances generated are computed using Laplace transforms, Fourier transforms, and Brigg's methods. The results show that the disturbances generated from a pointsource which is sinusoidal in the azimuthal direction are confined in a region 70–120 degrees as expected from the 2D eigenvalue problem.

In the future, we plan to compute the evolution of disturbances generated from isolated roughness elements.

Parallelization of the Jet Simulation Code

Ayodeji Demuren

The computer code used for direct numerical and large eddy simulation of complex 3D jets is to be parallelized to run on multiple processors. Scalability, portability, and accuracy are considered of paramount importance. The goal is to be able to run the code on any available multiprocessor system, such as workstation clusters, Beowulf systems, or SMP systems from CRAY, SGI, HP, or Sun.

The chosen approach was to convert a previously developed higher-order-accurate numerical formulation for the simulation of complex jets to enable parallel simulation via domain decomposition. MPI library calls are used for communication. Stability of the multi-domain simulation is assured by use of difference operators which satisfy summation by parts criteria. Parallelization is completed for the convection-diffusion sections of the code and has been tested successfully on Beowulf, SMP systems, and workstation clusters. The Poisson solver part is currently being parallelized.

Future plans are to complete parallelization tasks and then optimize code for performance.

Simulation of Supersonic Jet Mixing in Lobe Ejectors

Chester E. Grosch

Mixing enhancement of high and low speed streams is utilized as a means to improve efficiency of supersonic combustors, reduce aircraft signatures, and control high speed jet noise. One common method of mixing enhancement is to use lobe mixer ejectors. Another is to place tabs on the edges of the jets. Only a few experimental studies are available to evaluate the performance and guide the design of these mixers. The objective of this research is to use numerical simulation to examine the performance of lobe ejectors, both with and without tabs, in order to understand the physics of the mixing and how it is affected by changes in the parameters of these devices.

A set of numerical calculations is carried out using the compressible, three-dimensional, time-dependent Navier-Stokes equations. Tabs are modeled by pairs of counter-rotating vortices. Various geometric configurations of the lobe mixers are simulated with periodic side boundary conditions to simulate an array of these devices.

The simulations of the lobe mixer without tabs shows that the jet becomes unstable and oscillates in the "sinuous" mode. For a particular lobe geometry and velocity ratio, the oscillation has a constant, narrow-band frequency near the inflow. Further downstream the amplitude grows and the motion becomes nonlinear leading to spectral broadening. Typical Strouhal numbers of the narrow band oscillation are about 0.45. As the disturbances become nonlinear, rapid mixing between the supersonic and subsonic jets occurs and by about half way down the channel, the jet and coflow become nearly fully mixed.

The origin of the oscillation may be the instability of the shear layers. The instability of a single compressible jet has been understood for over 30 years. However, nothing appears to be known about the instability of a periodic array of compressible jets. The effect of the periodicity on the instability is conjectured to be important, particularly as the spacing between the jets is small in a typical lobe geometry. Theoretical calculations of the stability of this configuration has been begun.

A set of simulations of the same geometry with tabs was completed. The results of the first of these has been partially analyzed. A final set of simulations of a pair of jets with solid side walls has begun.

Flow Diagnostics Using Laser-induced Thermal Acoustics

Roger Hart

The non-intrusive optical measurement of gas-phase parameters such as temperature, flow velocity, and pressure is of considerable utility in understanding the airflow around a test body in a wind tunnel. Laser-induced thermal acoustics (LITA) is a relatively new optical diagnostic method that has great promise for becoming a practical, accurate flow characterization tool. Two lasers are employed in LITA. The short pulse from the first laser creates a pair of counterpropagating acoustic wavepackets. A long pulse from the second laser is diffracted by the wavepackets onto a detector. Analysis of the various features of the resulting LITA waveform allows the determination of the speed of sound in the medium (and thus the bulk temperature), one or more components of the flow velocity, and the density or pressure. Advantages of LITA as compared to other, better-developed diagnostics are: LITA allows seedless velocimetry; LITA measurements take only about 1 microsecond, giving the potential for very high repetition rates for the study of turbulent flows; and LITA gives excellent ($\tilde{1}$%) single-shot accuracy and precision. The goals of the current work are: 1) to completely understand the physics of the LITA measurement process; 2) to express that understanding in a quantitative model which has been carefully validated against laboratory experiment; and 3) to design and test a simple, rugged, accurate LITA measurement system.

The fundamental optical and acoustical mechanisms of LITA are well understood; nevertheless, combining these to create a model that can accurately and robustly duplicate the results of well-controlled experiments has involved considerable effort. Several surprising insights into the best way to do LITA have resulted from this work. Experimental efforts to verify these results are currently under way. The bulk of recent effort has centered on designing a LITA apparatus using diffractive optics and an improved homodyne detection scheme. It is anticipated that this approach will be both more robust and more accurate.

We are currently scheduled to make LITA flow measurements in NASA Langley's Basic Aerodynamic Research Tunnel (BART) in the late summer of 2000. Major effort during the next two reporting periods will be devoted to designing and testing instrumentation for that test. This research was conducted in collaboration with R.J. Balla and G.C. Herring (NASA Langley).

Generation of Nonlinear Difference Equation Models from Data for Flow Control

Richard Longman

Modern feedback control system design methods usually make use of a system model as part of the control action computed in real time, and these methods apply to ordinary differential equation models and to difference equation models. Although there is a field of control of systems governed by partial differential equations, it is not well enough developed to produce results easily in practical problems. Hence, from the point of view of making practical controller designs for flow control, it is best to avoid partial differential equation models. But even then, ordinary differential equation or difference equation models of fluid flow produced, for example, from Navier-Stokes equations, are far too complex and too hard to generate to be of use in control design. This research is aimed at starting from data to produce a fit using nonlinear difference equation models, so as to produce a model that is as simple as possible to model the effects of importance in the situation of interest. We are currently using a flow solver to create simulated data, and then doing numerical experiments for producing fits of the data with nonlinear difference equation models, or nonlinear ARX models. The simulated data used at the moment is for a plunging airfoil in transonic flow. The input is the input velocity of the airfoil, and the output quantity of interest is the lift. The difference equation's job is limited to producing this single scalar quantity, the lift occuring in a specialized situation, and hence it can potentially be quite simple, much simpler than a complete flow solver that creates a complete description of everything about the flow. In producing linear difference equation models one must make various decisions. At one extreme is a discrete nonlinear Volterra series, which represents the output as a linear combination of past inputs and input moments. Another possibility is to use moments of not only the input, but also the output, producing a nonlinear ARX model. In either case one must pick the appropriate number of terms, and the appropriate number of terms for input quantities and for output quanties. We have developed some theoretical understanding of the tradeoffs between these model types. We are currently conducting numerical experiments to establish ways of picking the order needed, and to establish how one makes the decisions of what type of model to use and what nonlinear terms need to be included. Two sets of data are always used, one for making the fit, and an independent one for testing the resulting model. We are aiming to develop techniques to determine when to stop adding terms that improve the fit to the fit data set, but do not assist in improving the prediction ability on the second data set..

This research was conducted in collaboration with Brian Allan (ICASE) and Jer-Nan Juang (NASA Langley). Brian Allan brings the needed expertise in flow modeling and CFD to the problem.

High Temperature Piezoelectric Polyimides

Zoubeida Ounaies

This work is motivated by NASA's interest in developing high performance high temperature piezoelectric polymers. Towards that end, a number of polyimide polymers were synthesized and investigated for this particular study. Currently semicrystalline PVDF and its copolymers are the only commercially available piezoelectric polymers. However, this family of vinyl polymers does not possess the thermal stability dictated by our high temperature aerospace applications. Polyimide polymers, on the other hand, have unique and attractive features for electronic applications. Due to their exceptional thermal, mechanical, and dielectric properties, polyimides are already widely utilized as matrix materials in aircraft and as dielectric materials in microelectronic devices. Particularly interesting is the potential use of piezoelectric polyimides in micro electro mechanical systems (MEMS) devices since piezoelectric vinyl polymers such as polyvinylidene fluoride (PVDF) do not possess the chemical resistance or thermal stability necessary to withstand conventional MEMS processing. The intent of this project is to elucidate the mechanism and key components required for developing piezoelectricity in amorphous polymers and further to apply this understanding in designing a novel high temperature piezoelectric polyimide.

Early molecular modeling indicated the potential of forcing piezoelectricity in amorphous polyimides by adding dipole functionalities to the polymer structure. This was the impetus for synthesizing a series of polyimides to investigate their potential piezoelectric response resulting from adding various dipoles to the structure. The dipoles can be either pendant to the polymer backbone or reside within the main or within a side group of the polymer. The effect of structural changes, including variations in the nature and concentration of dipolar groups, on the remanent polarization and piezoelectric coefficient is examined. Fundamental structure-piezoelectric property insight will enable the molecular design of polymers possessing distinct improvements over state-of-the-art piezoelectric polymers including enhanced polarization, polarization stability at elevated temperatures, and improved processability.

Future work will explore the feasibility of developing new monomers and polymers with larger dipole concentrations while maintaining chain flexibility and favorable geometries.

This research was conducted in collaboration with Joycelyn S. Harrison (NASA Langley).

Piezoelectric Ceramics for Use as Actuators

Zoubeida Ounaies

The intent of this project is to assess the ability of piezoelectric wafers to meet design requirements in aeronautical applications (such as active control of aircraft), and to improve effective piezoelectric properties by modifying processing and geometry of actuators. The immediate goal is the characterization of THUNDER (pre-stressed piezoelectric unimorph) actuators to gain a fundamental understanding of how processing variables affect actuator performance.

Measurements completed so far include dielectric properties, as well as displacement and force as a function of boundary conditions and frequency. The capacitive behavior and power consumption of the piezoactuators is also determined. Empirical relations are developed from experimental data to predict the capacitance and loss tangent of a piezoelectric ceramic as nonlinear functions of both applied peak voltage and driving frequency. It was demonstrated that by incorporating the variation of capacitance and power loss with voltage and frequency, satisfactory estimates of power requirements can be obtained. These relations allow general guidelines in selection and application of piezoelectric actuators and driving electronics for active control applications.

Future work will include determining the power efficiency and the electro-mechanical conversion coefficient of the pre-stressed actuator, and relating the electronic resonance behavior to specific geometrical variables.

This research was conducted in collaboration with Thomas L. Jordan (NASA Langley).

A Hysteresis Model for Piezoelectric Ceramics

Zoubeida Ounaies

This project involves characterizing the hysteresis behavior of piezoelectric ceramics, both computationally and experimentally, under moderate to high fields in order to estimate their nonlinear behavior. This can be useful in the implementation of control design where the prediction of the behavior of the actuator throughout the full range of operation is necessary.

Experimental validation of the hysteresis model for lead zirconate titanate ceramic (PZT) was successfully accomplished. Three representative PZT compositions were considered namely PZT4, PZT5A and PZT5H. A family of hysteresis loops were measured for all three compositions, and excellent agreement between the data and the model was attained.

Future work involves extending the model to include frequency and temperature effects for the above mentioned PZT compositions. We are also interested in understanding the effect of pre-stressing inherent to some actuators (such as THUNDER and RAINBOWS) on the nonlinear behavior as a function of applied field and frequency.

This research was conducted in collaboration with Ralph C. Smith (North Carolina State University).

Computation of Three-dimensional Acoustic Fields

Alex Povitsky

Compact numerical schemes are widely used for challenging problems of computational physics including aeroacoustics. The number of arithmetic operations per grid node is approximately equal for explicit and compact formulations of the same order. Poor computational efficiency of compact numerical schemes in comparison to explicit numerical schemes is explained by architectural features of modern computers, namely, use of cache memories and use of multi-processor computers with distributed memory. It is known from computational experiments that the aeroacoustic propagation problems for nonuniform mean flows suffer from exponentially growing instabilities which actually are important noise sources. For instance, this situation holds near a stagnation point of a mean flow such as the leading edge of an airfoil. However, in available studies the complicated behavior of aeroacoustic fields near stagnation points have not received special attention. The objective of this research is to propose high-order numerical methodology suitable for modern computers and to use this methodology to understand physical features of noise generation and propagation.

We have successfully demonstrated efficient parallelization of a multi-dimensional aeroacoustics problem solved by a compact scheme. To achieve good parallelization efficiency, the schedule-driven computations are performed in such a way that an idle stage for the processors is avoided and the number of inter-processor messages is minimized. The experimentally obtained parallelization efficiency is comparable to that for the explicit dispersion-relation-preserving scheme (DRP) with the same order of accuracy. The parallelization results for the DRP scheme are provided by group of Philip J. Morris (Pennsylvania State University). To reduce the main memory-cache exchange for compact schemes, the cache-aware compact numerical algorithm has been developed theoretically. This algorithm is designed in such a way that the data pass through cache only twice for both 2D and 3D cases and for any number of cache levels. Now that the methods of efficient implementation of compact schemes on non-von Neumann computers are understood, we can investigate details of noise propagation for nonuniform mean flows.

We are working on numerical simulation of the aeroacoustic pulse dynamics in an incompressible stagnation flow. We plan to investigate efficiency of the proposed cache-friendly algorithm with help of specialized tools for the main memory-cache exchange measurements.

Shock Wave Propagation in Weakly Ionized Gases

Robert Rubinstein

It has been proposed that electrostatic forces can make large contributions to the pressure in weakly ionized gases; since the Coulomb force is long-range, this additional pressure could increase the sound speed in weakly ionized gases substantially. This explanation of the reported 'anomalous effects' in shock waves in weakly ionized gases is particularly promising because it predicts a large sound speed increase even if the ionization level is very low. The present work assesses this proposal by numerical simulation and theoretical analysis.

Calculations of shocks in weakly ionized gases show that a charge separation region exists ahead of the shock, provided that the Debye length is larger than the neutral shock thickness. The proposed effect was modeled numerically by adding a term proportional to the charge difference to the neutral pressure. The simulations show that the shock can be accelerated, despite the very small lengths over which appreciable charge separation exists. These results then raise the question of how the neutral pressure can be enhanced by the existence of unbalanced charge. Since the ion component in the charge separation region is a strongly non-ideal gas, a theoretical model was formulated in which an ideal gas is collisionally coupled to an Enskog non-ideal gas. In this coupled system, the pressure in the ideal gas can be shown to be increased by the presence of the non-ideal gas. But in the actual neutral-ion system, the Coulomb forces do not act on the neutrals. This means that the couplings in the theoretical model are modified in such a way that the increased pressure in the ideal gas disappears. We conclude that the proposed mechanism does not lead to the required effect.

This research was conducted in collaboration with A.H. Auslender (NASA Langley).

Multiple Scale Turbulence Modeling

Robert Rubinstein

In turbulent flows which arise in practice, different dynamic processes can be dominant at different scales of motion. Such flows, which include strongly time-dependent flows and flows subject to large external agencies like buoyancy and curvature, cannot be described by a single characteristic length scale. An alternative to current models which characterize turbulence by a single length is the multiple scale model, in which the length scale determining equation is replaced by a system of equations. Recent research on this class of models led to the development of a model in which spectral transport and dissipation rate are distinct and satisfy separate equations. A defect of this model is the strong Reynolds number dependence of the coefficients in the dissipation rate equation which makes numerical implementation very difficult.

The model was reformulated by allowing the spectral transport to take different values in different spectral regions. Simple closures for the energy transfer through each region were proposed. The transport into the region of smaller scales then equals the dissipation rate. The coefficients in the resulting equations are all independent of Reynolds number.

Multiple scale modeling may also prove useful in connection with modeling turbulent diffusion. Development of closures for the diffusion terms in the transport equations based on multiple scale ideas is in progress.

This research was conducted in collaboration with K. Hanjalic (Technical University of Delft, Netherlands).

Theory of Rotating and Stratified Turbulence

Robert Rubinstein

Rotating turbulence can be modeled statistically as the turbulence of weakly interacting inertial waves. In this theory, energy transfer is due to interactions between sets of three resonant wave modes. If these triad interactions can be suppressed, then four-wave interactions become dominant and new dynamic possibilities arise, including a three-dimensional inverse cascade to large scales. In small aspect ratio domains, three wave resonances at the largest scales of motion are not possible. However, attempts to generate the three-dimensional inverse cascade in numerical simulations of forced rotating turbulence in small aspect ratio domains have not been successful. The possibility that the inverse cascade could be important in geophysical applications motivates the present investigation of the reasons for this failure.

In small aspect ratio domains, resonant triads exist in which all three modes have the same magnitude as the inverse domain height. It is proposed that in the simulations with isotropic forcing at scales smaller than the inverse height, these triads transfer energy to the perpendicular plane, rapidly causing the largest scales of motion to become two-dimensional. If the force is anisotropic, such triads do not exist. Simulations in a cubical domain with a strongly anisotropic force demonstrate that a range of three-dimensional large scales exist. However, it cannot yet be concluded that these scales are caused by an inverse cascade since the very largest scales are again two-dimensional.

Simulations are now in progress in a small aspect ratio domain with anisotropic forcing. This problem should be free of resonant triads: it is expected that if rotating turbulence dominated by four-wave resonances can exist at all, it must exist in this problem.

This research was conducted in collaboration with L.M. Smith (University of Wisconsin). Extensive correspondence with C. Cambon (Ecole Central de Lyon, France) is also acknowledged.

Turbulence Model for Rotating Flows Based on the Energy Spectrum

Ye Zhou

Rotating turbulent flows are quite common in engineering practice as well as a wide range of practical applications of relevance to geophysical and astrophysical systems. The flow is also of fundamental interest from a turbulence modeling point of view since the influence of rotation on turbulence is quite complex. The formulation of turbulence models to accurately describe the essential features of rotating flows requires that the energy transfer mechanisms are effectively incorporated in the governing transport equations.

A generalized eddy viscosity model is formulated by using the rotation modified energy spectrum. Rotation and mean shear effects are directly included in the eddy viscosity without the use of the local equilibrium assumption. The formulation also includes the modeling of vortex stretching and viscosity destruction terms of the dissipation rate equation based on the limit of rotating isotropic turbulence at high Reynolds numbers. The model is shown to reproduce the dominant effects of rotation in rotating homogeneous shear flows and turbulence channel flows subject to span-wise rotation.

We plan to apply this model to other turbulent flows where rotation effects are important.

This research was conducted in collaboration with S. Thangam and X.-H. Wang.

COMPUTER SCIENCE

Object-oriented Parallel Load Balancing Library

Abdelkader Baggag

The aim of this project is to build a general purpose parallel load balancing library for use with Finite Element packages. The library will initially be used in conjunction with an Aeroacoustics Solver based on the Discontinuous Galerkin method, but it is being written in an object-oriented language so that its functionality may be extended for use with other applications without major changes to the overall structure.

Most load balancing packages require the user to extract the necessary information from their data structures, which can be a non-trivial task. By using an object-oriented language like C++, we can write an interface which is easier and more natural to use. As the library is intended for use with Finite Element packages, the user will be required to pass in the information in terms of the elements in the grid and the library will then extract the required details to build the graph. The actual load balancing algorithm we are currently considering is similar to a diffusive routine used by Linda Stals in her Finite Element code; other methods shall be added later on. Work on this project only started recently and we are still working on the design of the various components; however, we have started implementing the diffusive load balancing algorithm.

We plan to continue testing the current implementation and to add other load balancing paradigms to the library, such as a multilevel diffusion algorithm.

This research was conducted in collaboration with Linda Stals (ICASE and Old Dominion University).

Parallel Implementation of the Discontinuous Galerkin Method

Abdelkader Baggag

Unstructured grids are readily generated for geometrically complex problems; however, high-order finite-difference and finite-volume methods commonly used by the CFD community are not robust when applied on such grids. Discontinuous Galerkin is a robust and compact method that provides a framework for the development of high-order methods for unstructured grids. Such a method has been developed by H.L. Atkins (NASA Langley), and applied to the time-dependent problem of acoustic scattering. The long-term goal is to develop methodology to autonomously and locally adapt the grid, the order of approximation, and even the choice of the governing equations. Atkins' original code was sequential, object-oriented, and cleanly modular. Although the basic code is easily and efficiently ported to parallel platforms, the desired adaptation features pose problems for parallelization and load balancing and must be addressed.

Several parallelization approaches have been studied and evaluated. A more natural and symmetric approach has been chosen and implemented in an object-oriented code. The parallel implementation is MPI-based and has been tested on various parallel platforms such as the SGI Origin, IBM SP2, and clusters of SGI and Sun workstations. The scalability results show slightly superlinear speedup on fixed-size problems due to cache effects.

A careful performance evaluation will be undertaken, including modeling and experimentation. A robust load balancing methodology that is independent of the source of the imbalance is being developed. This is required to facilitate the use of h-p adaptive methods as well as adaptive physics (different equations solved in different regions).

This research was conducted in collaboration with Harold Atkins (NASA Langley) and David Keyes (ICASE and Old Dominion University).

Application of Parallel and Distributed Computing to Visualization and Data Fusion Problems in the Digital Earth

Thomas W. Crockett

To implement the Vice President's vision of a Digital Earth, vast quantities of data from disparate sources must be integrated into an intuitive, accessible representation. A fundamental building block in such a system is a high-resolution, interactive model of the physical Earth.

We are investigating the use of parallel and distributed computing technologies (processing, I/O, networking, etc.) to provide interactive data fusion and visualization services for Digital Earth applications. Our initial focus is the development of a parallel visualization tool for rendering high-resolution, high-quality images of the Earth's surface by combining elevation data with satellite imagery. We have demonstrated this tool (called EarthView) on Langley's SGI Origin2000 system and ICASE's PC-based Coral cluster. We are currently adding support for triangle strips to the underlying parallel renderer used by EarthView. Triangle strips are ideally suited for rendering of high-resolution latitude-longitude grids, and we expect this capability to reduce memory requirements and improve performance by factors of 2–3, providing interactive frame rates for scenes containing 10–20 million grid points.

In the near term, we plan to extend EarthView for use with higher resolution data, such as USGS's 30-arc-second global elevation dataset (933 million grid points). This will require the development of parallel multi-resolution capabilities and improved clipping techniques, along with parallel I/O for rapid access to data which is cached on disk. Ultimately, the EarthView application is intended as a substrate on which other types of data can be combined and visualized. We are especially interested in atmospheric data from sources such as Langley's LITE experiment and the upcoming Picasso/CENA mission. We believe that the techniques developed here will lead us toward Digital Earth's goal of providing interactive access to multi-petabyte datasets, a challenge which is beyond the capability of current computing technology.

Scalable Parallel Volume Rendering for Large-scale Unstructured Grids

Thomas W. Crockett

Three-dimensional unstructured grids play an increasingly important role in a variety of computational tasks, including aerodynamics simulations of interest to NASA. Since these grids are often used in conjunction with complex vehicle geometries to simulate complex physical phenomena, the ability to visualize the data is often critical in understanding the results. Among the visualization techniques available for 3D data, volume rendering is unique in its ability to convey information throughout the depth of a domain. However, volume rendering is a computationally demanding task, especially for large unstructured datasets.

In order to boost the performance of volume rendering methods and to apply them to state-of-the-art computations, we have been developing parallel volume rendering algorithms for unstructured-grid data. Our current algorithm was originally developed and tested on the IBM SP2 and Intel Paragon for datasets with roughly 0.5 million cells, using up to 256 processors. Since then we have refined the algorithm, and have recently extended our results to a much larger CFD dataset containing 18.2 million cells. Our results on the Cray T3E at Goddard Space Flight Center have been encouraging, showing that communication and other overheads remain manageable with as many as 512 processors. Rendering times range from 5–20 seconds per frame at typical image resolutions. We have also ported our code to ICASE's 32-node PC cluster, where it exceeds the T3E's performance at comparable numbers of processors. These results demonstrate the feasibility of volume rendering large unstructured-grid datasets at near-interactive rates on a variety of parallel architectures.

In order to obtain optimum rendering performance, several algorithmic parameters need to be tuned for both the architecture and the dataset. Determining these values is not straightforward; in fact, we have identified nearly a dozen variables that influence performance of the renderer. For efficiency and ease of use, we would like to predict what settings will optimize performance for any given application. We therefore plan to refine our analytical performance models in order to obtain a better understanding of the complex relationships between various parameters. The ultimate goal would be to incorporate an automatic mechanism to dynamically tune the algorithm based on characteristics of the dataset and the architecture, as well as viewing and visualization parameters.

This research was conducted in collaboration with K.-L. Ma (ICASE).

Parallel Implicit Solvers for Simulation of Multiscale Phenomena
David E. Keyes

The development and application of parallel implicit solvers for multiscale phenomena governed by PDEs are our chief objectives. Newton-Krylov-Schwarz (NKS) methods have proven to be broadly applicable, architecturally versatile, and tunable for high performance on today's high-end commercial parallel platforms (e.g., Cray T3E, SGI Origin, IBM SP). Both structured-grid and unstructured-grid CFD legacy codes have been ported to such platforms and reasonable objectives for algorithmic convergence rate, parallel efficiency, and raw floating point performance. However, architectural challenges have increased on the next generation of high-end machines, as represented, for instance, by the ASCI machines at Lawrence Livermore, Los Alamos, and Sandia National Laboratories, and also on Beowulf clusters, such as ICASE's Coral. Our primary efforts are concentrated on algorithmic adaptations of NKS methodology appropriate for the emerging architectures and on evaluation of new software tools and methodology to get the most performance out of them.

The general approach embodied in the NKS family of algorithms is documented in previous ICASE technical reports, among other places. We solve steady-state conservation laws with pseudo-timestepping as a continuation technique, using Newton's method on each timestep, truncated GMRES to approximate the solution of the Newton correction equations, and parallel Schwarz preconditioning of GMRES. Specific emphases in the most recent reporting period include enhanced understanding (modeling and instrumentation) of per-node floating point performance, introduction of operator-adaptive preconditioning of alternating block factorization (ABF) type, and further porting and scalability analyses.

There are two major phases in any implicit solver for discrete conservation laws, including NKS: flux evaluation and algebraic solution (with a sparse Jacobian). Using the STREAM benchmarks to define memory bandwidth limits and instruction counting in compiled code, we have determined that the flux evaluation phase is generally load-limited and the linear algebra phase is generally bandwidth-limited. For neither phase is the peak performance rating of the hardware (clock frequency times superscalarity) a useful predictor of performance, but actual performance in the respective phases follows the limiting models above.

We have employed an alternating block factorization preconditioner in place of block-ILU (BILU) on each subdomain of the Schwarz preconditioner and shown an improvement over BILU throughout a range of parameterized radiation diffusion problems. At one extreme diffusion dominates over intercomponent coupling, at the other intercomponent coupling dominates over diffusion. Alternating block exploits the simpler structure of each of these individual phases to construct a composite preconditioner. Multilevel methods may be used with greater ease and advantage on the decoupled elliptic problems in the diffusively-dominant phase than in the composite problem.

Finally, on the ASCI "Red" multiprocessor at Sandia, we have run on up to 2,048 nodes and achieved 0.156 Teraflop/s in the attempt — a new record for computational rate on a (meaningful) PDE-based computation, over the 0.111 Teraflop/s achieved on a structured PDE computation on the Japanese numerical wind tunnel (NWT) vector multiprocessor in 1996. This computation is a Bell Prize finalist entry at Supercomputing'99. (The 1997 and 1998 winning entrants were not based on PDE computations.)

We will continue to develop NKS methods in implicit parallel CFD, examining a variety of algorithmic, programming paradigm, and architectural issues. We will seek to exploit the existence of a highly optimized parallel PDE solver in PDE-constrained optimization problems. We will also increase the complexity of the models in our NKS radiation transport work, in accordance with the ASCI project roadmap.

This research was conducted in collaboration with W. Kyle Anderson (NASA Langley), Dana Knoll (Los Alamos National Lab), Dinesh Kaushik and Xin He (Old Dominion University), and Satish Balay, William D. Gropp, Lois C. McInnes, and Barry F. Smith (Argonne National Laboratory).

Temporal Process Logic

Gerald Lüttgen

This work is inspired by NASA's search for a suitable ***design methodology*** for ***mode logics*** of ***flight guidance systems.*** The methodology should support the transitioning between ***system requirements*** and ***system design*** and should also allow for the efficient analysis of mode logics. It has previously been shown that ***model checking*** is an efficient technique for reasoning about mode logics designs that are ***operationally*** specified by ***state machines*** and where system properties are ***assertionally*** stated as ***temporal logic formulas.*** The objective of the present research is to develop a framework which allows one to mix state machines specifications and temporal logic formulas. As indicated, the framework should support a notion of ***stepwise*** refinement which is consistent with model checking, i.e., the refinement relation must be compatible with the gradual replacement of temporal formulas by state machines satisfying the formulas.

The key for approaching this task is the observation that ***Büchi*** automata may serve as a uniform framework for integrating operational and assertional styles of specification, since properties in ***linear-time logic*** (LTL) can be translated to Büchi automata in a way that model checking reduces to a language inclusion problem. Moreover, we found that DeNicola and Hennessy's theory of ***must-testing*** for state machines provides an adequate basis for defining the desired refinement relation. In this theory, a state machine must pass a ***test***, given by a state machine with acceptance condition, if every run that both state machines can do in parallel satisfies the test's acceptance condition. We extended must-testing to Büchi automata such that a state machine satisfies an LTL formula, represented as a ***purely nondeterministic*** Büchi automata, if and only if for all tests the following holds: if the formula must pass the test then also the state machine must pass the test. Our theoretical results guarantee that system properties are preserved when gradually transitioning from system specifications to system designs. Thus, our notion of must-testing for Büchi automata provides a sound and elegant semantic basis for mixed specification/design languages.

Future work includes the definition of a language which combines assertional and operational styles of specification and whose semantics can be given in terms of Büchi automata. It should also be investigated how the methodology of stepwise refinement, as described above, can be integrated in popular software design tools which are expected to be used for the development of future flight guidance systems.

This research was conducted in collaboration with Rance Cleaveland (SUNY at Stony Brook).

Statecharts, Compositionality, and Intuitionistic Logic

Gerald Lüttgen

Statecharts is a popular visual design notation for embedded systems, which extends state machines by the concepts of concurrency, hierarchy, and priority. Unfortunately, most Statecharts variants do not possess a *compositional* semantics. Compositionality, however, is a necessary prerequisite for the composability of specifications and for the re-usability of behavioral properties. The objective of the present work is to develop a compositional and *fully-abstract* semantics for Statecharts.

We analyze the semantic properties of Statecharts by employing ideas and techniques from *intuitionistic logics*. Traditionally, Statecharts are mapped to ordinary state machines where transitions are referred to as macro steps. Each macro step results from a sequence of micro steps which are enabled in a given Statecharts configuration. Our approach to the compositionality problem encodes a Statecharts configuration by an intuitionistic formula. We interpret these formulas in specialized intuitionistic Kripke structures, namely linear chains. Thereby, we capture the causality inherent in a Statecharts macro step in an operational and compositional fashion. The logical approach allows us to derive a number of important results, such as a full-abstraction theorem for our Statecharts semantics with respect to the traditional step semantics of Pnueli and Shalev.

In the future, we intend to complete the intuitionistic framework by adding an algebraic theory for macro-step semantics. Our results can then be used to derive a structural operational semantics that provides a sound and compositional basis for implementing Statecharts.

This research was conducted in collaboration with Michael Mendler (Sheffield University, England).

State-space Generation Using Multi-valued Decision Diagrams

Gerald Lüttgen

Many formal verification techniques rely on the construction of the *state space* of the system under consideration. For coping with the large complexity of systems in practice, efficient data structures, e.g., *decision diagrams*, have been developed which encode state spaces in sub-linear space. However, while this approach exhibits excellent results compared to traditional techniques, it can still be problematic for realistic systems. In fact, the use of decision diagrams may make the state-space generation a *time-bound* problem instead of a *memory-bound* problem. The objective of this work is to speed up the state-space generation of *asynchronous systems* by exploiting the *locality of system events*.

We approach the task by defining a state-space generation algorithm that works on *multi-valued decision diagrams* (MDDs) that are 'horizontally' *partitioned into levels*. The locality of events in asynchronous systems can then be exploited to optimize existing MDD-based algorithms regarding their time efficiency. More precisely, most events of an asynchronous system only affect a few local components of the global state, such that in many cases operations on MDDs can be performed at a few levels (often at just a single level), rather than having to traverse the entire data structure. Although this approach involves some space overhead due to additional bookkeeping and also requires some sophisticated techniques for consistently *handling caches*, it is considerably more time-efficient than existing MDD-based algorithms. We verified this statement by prototypically implementing different variants of our algorithm within the *Petri net tool* SMART and by measuring their performance when being applied to standard example systems taken from the literature. Although our algorithms are not optimized, they are on average significantly faster (measured speed-ups ranged between 20% and 2000%) than related algorithms that do not exploit locality.

In the future, we want to further optimize our MDD algorithms and their implementations. One important issue that needs to be addressed concerns the efficient handling of caches. As soon as our implementation is fine-tuned, we will apply verification techniques based on MDDs to analyzing systems of relevance to NASA Langley. Moreover, our algorithm seems to be suitable for an efficient parallel implementation on distributed computing environments, such as ICASE's Beowulf system, Coral.

This research was conducted in collaboration with Gianfranco Ciardo and Radu Siminiceanu (The College of William & Mary).

Arcade: A Distributed Computing Environment for ICASE

Piyush Mehrotra

Distributed heterogeneous computing is being increasingly applied to a variety of large-size computational problems. Such computations, for example, the multidisciplinary design optimization of an aircraft, generally consists of multiple heterogeneous modules interacting with each other to solve the problem at hand. These applications are generally developed by a team in which each discipline is the responsibility of experts in the field. The objective of this project is to develop a GUI-based environment which supports the multi-user design of such applications and their execution and monitoring in a heterogeneous environment consisting of a network of workstations, specialized machines, and parallel architectures.

We have been implementing a Java-based three-tier prototype system which supports a thin client interface for the design and execution of multi-module codes. The middle tier consists of logic to process the user input and also to manage the resource controllers which comprise the third tier. Among other enhancements to our implementations, we have been focusing on resource management and monitoring, and module execution across multiple domains. We have been investigating various methodologies to provide secure mechanisms for remote executions of modules. This includes not only the actual running of the code but also the transfer of the input data and the results in a secure manner across domains. We have implemented several techniques using different technologies, e.g., ftp and RMI and are currently in the process of comparing the approaches. For resource monitoring and management, we have focussed on Jini, a new technology from Sun. Jini uses multicasting for its *discovery* and *join* protocols and hence works well with a single domain which supports multicasting. We have extended Jini with a hierarchical *tunneling service*, so that Jini services can be used across non-multi-castable networks. The implementation has been tested across the ICASE and ODU domains and is being incorporated in the Arcade environment.

In the future we will continue enhancements of the whole environment. In particular we will focus on developing an XML-based resource description system so as to allow the resource management system to support more dynamic features. We would also like to replace our current scripting language with a more powerful and standardized approach, such as Python or an XML-based system.

This research was being conducted in collaboration with A. Al-Theneyan and M. Zubair (Old Dominion University).

Languages for High Performance and Distributed Computing

Piyush Mehrotra

There are many approaches to exploiting the power of parallel and distributed computers. Under this project, our focus is to evaluate these different approaches, proposing extensions and new compilation techniques where appropriate.

OpenMP is a set of directives extending C, C++, and Fortran that provides a shared-memory parallel programming model. Current parallel architectures are built by interconnecting nodes that internally provide true-shared memory across a small number of processors. Both hardware and software-based approaches are used to provide a shared address space across the physically distributed memories of the nodes in larger systems. OpenMP provides an easy and incremental approach for small-scale shared memory systems. However, controlling and exploiting data locality becomes an issue as the latency for data transfers across the nodes increases for larger systems. Based on our experience with HPF, we had earlier proposed a set of OpenMP extensions which allows users to control the locality of the data and specify the distributions across the nodes. Our experience with codes has shown that our current set of extensions is not enough. In particular, users should be able to control the distribution of data to the processors within each node. Also, the layout of the data within each node can also become a performance issue. We are currently investigating alternatives which would allow such support.

In the future, we will continue investigating the efficacy of the directives, both from the view of performance and user specification.

This research was conducted in collaboration with B. Chapman (University of Houston).

Towards the Verification of Java Byte-code

César Muñoz

Java has emerged in the last years as the programming language for the Internet. Most of Java's success is due to the emphasis the designers of the language were giving to portability. In fact, a key aspect of the language is the specification of a standard execution model, called the Java Virtual Machine (JVM), which executes machine-independent byte-code. Thus, when a Java application is running in a heterogeneous network, the byte-code is transfered between the nodes, and locally executed by the JVM of each particular node. Security and safety concerns on the execution of the byte-code are particularly important in embedded systems like vehicle navigation or air traffic control applications, where reliability is a critical issue. Byte-code validations are performed by the JVM mainly through a type-checking algorithm. However, holes in the security policy for Java byte-code have been discovered. In this work we apply Formal Methods to study Java byte-code verification in order to improve the safety aspects of the language.

Model-checking, usually based on state-space exploration, is a mature technology for automatic verification of finite-state systems. To handle the case of infinite-state systems, as the case of Java applications is in general, techniques involving abstraction and theorem-proving are necessary. This work focuses on a reasonable subset of the Java byte-code language containing threads, and the basic types integers and booleans. For this subset of the language, an abstraction technique which uses enriched type information on execution points of a byte-code program will be developed. Theorem proving will be necessary to build the appropriate abstraction, and model-checking could be used to automatically discharge properties on finite-state abstractions. So far, we have defined a formal model for Java byte-code. We have applied it to examples which show that existing techniques can be improved.

Once our initial formal framework has been completed, it will be extended to handle aspects not considered in the first approach, such as exceptions, and a richer set of instructions of the Java byte-code.

This research was conducted in collaboration with Peter Habermehl (University of Paris 7).

A Generic Higher-order Unification Algorithm

César Muñoz

Higher-order Unification (HOU) is the process to solve equations in the λ-calculus. In contrast to first-order unification, HOU is undecidable and a most general solution does not always exist. In dependent type systems, where terms are allowed to appear inside types, HOU reveals some additional technical difficulties. HOU problems arise in several program analysis techniques, for instance type checking, abstraction, optimization, and program transformation. HOU is also a main issue on advanced theorem provers and logic programming tools. Fundamental improvements in HOU algorithms are essential for the advancement of these efforts. In the past years, diverse techniques for HOU have been proposed. The objective of this work is to investigate the application of a new promising technique called ***explicit substitutions***, which allows substitutions to be first-class objects in λ-calculus, to implement a generic and efficient HOU algorithm for dependent-type systems.

It is well known that the λ-calculus can be encoded in a first-order setting via a nameless notation for variables called *de Bruijn indices*. Recently, the HOU problem for the simply typed λ-calculus has been considered in a suitable first-order equational theory via de Bruijn notation and explicit substitutions. First-class substitutions and de Bruijn indices notation give to the HOU problem a first-order presentation that eliminates technical problems related to the higher-order aspects of the λ-calculus. Furthermore, they allow a finer grained control over the unification algorithm. We have implemented, in the functional object-oriented language Objective Caml, a HOU algorithm for the Calculus of Constructions, the richest dependent-type theory. The implementation has been tested with some known theorems. Remarkably it has synthesized in a few seconds, given some initial logical simplifications, an automatic proof of a famous Cantor theorem. We have begun the meta-theoretical study of the algorithm.

We plan to extend the HOU algorithm in such a way that it can be used as a constraint solver module for theorem provers and logic programming tools.

This research was conducted in collaboration with Nikolaj Bjørner (Kestrel Institute).

Verification of the AILS Alerting Algorithm

César Muñoz

The Airborne Information for Lateral Spacing (AILS) is a project being conducted at the NASA Langley Research Center. Its objective is to reduce traffic delays and increase airport efficiency by enabling approaches to closely spaced parallel runways in Instrument Meteorological Conditions. Langley researchers have developed an algorithm called the AILS Alerting Algorithm to provide situational awareness, alerting, and guidance when one of the two aircraft on parallel approaches intrudes on the other aircraft's air space. This algorithm analyzes the parallel aircraft states and makes time projections of possible intrusion scenarios. Based on these projections and risk criteria, the algorithm triggers a sequence of four caution and warning alerts. The objective of this work is to apply formal analysis and verification techniques to the AILS alerting algorithm to discover any possible errors, leading to unalerted near misses or collisions, that have not been detected during testing and simulation.

The AILS algorithm involves simultaneous physical mathematical calculations and a decision management procedure. The approach taken in this work is the combination of two different techniques to handle both aspects of the algorithm:

- Mathematical calculations involving physical characteristics will be validated first by hand and then using computer algebra tools.
- The algorithm will be formally modeled in a specification language and analyzed using model checking, theorem proving, or any other appropriate tool.

This work is in its first stage of development. A preliminary validation of an informal description of the algorithm reveals suspicious behaviors. Since a formal description of the algorithm is not available, the ongoing work focuses on the understanding of the overall environment of the alerting algorithm as it is assumed by the implementation.

Once the formal analysis of the algorithm has been completed, research will be necessary to understand how computer algebra tools can be interfaced with standard formal verification systems.

This research was conducted in collaboration with Victor Carreño (NASA Langley).

Cache-friendly Algorithms in Scientific Computing

Alex Pothen

Modern computer architectures achieve high performance by issuing and executing a number of instructions in parallel. However, since execution times of instructions improve at a faster rate than memory access times, in many irregular computations the bottleneck is the speed with which the data operands can be fed to the CPU. This bottleneck is alleviated by means of caches, small fast memories to which data can be moved from the larger and slower primary memory, and from which data can be moved to the CPU. We study how irregular computations (e.g., unstructured mesh or sparse matrix algorithms) can be organized so that most of the data accesses can be performed within the cache for high performance.

We have studied the cache performance of the kernel of an unstructured mesh code for solving Euler's equations using a cache simulation tool called FastCache. We have reordered the data accesses in this computation with the Cuthill-McKee, Sloan, Nested Dissection, and space-filling-curve reordering algorithms. The cache performance of irregular codes is greatly enhanced by these orderings. We have looked at the number of cache misses generated by each statement in the kernel to understand the reasons why the reordering algorithms improve cache performance. We conclude that the CM and Sloan reorderings improve temporal locality, while Nested Dissection improves spatial locality. Our findings lead to general principles for improving the cache performance of irregular computations.

We are currently working on modeling multiple levels of caches, and on modeling analytically the cache performance of many kernels of scientific computing algorithms.

This research was conducted in collaboration with Jinghua Fu (ICASE and Old Dominion University) and Dimitri Mavriplis (ICASE).

Fast Algorithms for Incomplete Factorization Preconditioners

Alex Pothen

The parallel computation of robust preconditioners is a priority for solving large systems of equations iteratively in several scientific computing problems. We are developing parallel algorithms and software that can compute incomplete factorization preconditioners for high level fill.

We have developed a structure theory based on paths in the adjacency graph of the matrix to identify the nonzero elements in incomplete factors. We assume that the adjacency graph can be partitioned into subgraphs of roughly equal sizes such that few edges are cut by the partition. We map the subgraphs to processors, form a subdomain interconnection graph, and order the subdomains so as to reduce global dependences. On each subdomain, we locally reorder the interior vertices before the boundary vertices. This reordering limits the fill that joins a subgraph on one processor to a subgraph on another, and enhances the concurrency in the computation. The preconditioner computation takes places in two phases: in the first phase, each processor computes the rows of the preconditioner corresponding to the interior vertices of their subdomains. In the second phase, the rows corresponding to the boundary nodes are computed. This approach can make use of level-based and numerical threshold based algorithms for computing preconditioners in parallel.

Our preliminary results on the SGI Origin and Coral, the ICASE Beowulf cluster, show efficiencies in the range of 75% on up to 16 processors. We are continuing to develop our parallel implementation.

This research was conducted in collaboration with David Hysom (ICASE and Old Dominion University).

Spindle: An Algorithmic Laboratory for Ordering Algorithms

Alex Pothen

We have developed an algorithmic laboratory for quickly prototyping promising algorithms and experimenting with a collection of algorithmic variants for several ordering problems. Among these are the fill reduction problem: Order the rows and columns of the coefficient matrix to reduce the fill in sparse Gaussian elimination (both complete and incomplete factorizations); and the sequencing problem: Given a set of elements, and pairs of elements that are related, order the elements such that related elements are numbered consecutively. We employ object-oriented design techniques (OOD) to make the laboratory flexible and easy to extend.

OOD manages complexity by means of decomposition and abstraction. We decompose our software into two main types of objects: structural objects corresponding to data structures, and algorithmic objects corresponding to algorithms. This design decouples data structures from algorithms, permitting a user to experiment with different algorithms and different data structures, and if necessary develop new algorithms and data structures. We have implemented seven variants from the family of minimum degree ordering algorithms using this design paradigm. Prior to our work, there was no single code that implemented all of these algorithms. Our implementation makes it possible for us to change ordering algorithms midstream while ordering a problem. We have found this to be of benefit, since a hybrid algorithm that employs the multiple minimum degree (MMD) algorithm and switches at later stages to the approximate minimum degree (AMD) algorithm can improve performance for problems where either algorithm has poor performance. We have also analyzed the complexity of the minimum-degree algorithm.

We are currently considering the extension of our methods to ordering unsymmetric problems.

This research was conducted in collaboration with Gary Kumfert (ICASE and Old Dominion University). He has joined the Center for Advanced Scientific Computing at Lawrence Livermore National Lab as a computer scientist.

Efficient Solution of Radiation Transport Equations

Linda Stals

Under certain physical assumptions, such as isotropic radiation, optically thick material and temperature equilibrium, radiation transport may be modeled by a system of three nonlinear equations. Due to strong nonlinearities and large jumps in the coefficients these equations are challenging to solve. One tool which may be used to meet this challenge is to write an efficient parallel solver, thus, allowing us to solve a given problem more quickly. However, just considering the parallel efficiency is not enough, if we do not also take algorithmic scalability into account we cannot include details in the models which are of interest to the scientists. That is, the time required to solve the problem on a single processor should increase only linearly as the number of grid points is increased.

One technique we have explored is the use of adaptive refinement techniques to give more accurate answers for less computational work. Although more 'quantitative' analysis is needed, our initial results imply that our method is robust and captures the required information.

Radiation transport equations are time dependent. We have extended our code to include the time dependence and obtained some results for the test case where all temperatures are in equilibrium and the system of equations reduces to a single equation.

The initial phase of the time stepping algorithm is, computationally, the most difficult and the solution algorithms will not converge if the step size is too large. One of the aspects we would next like to focus on is the use of adaptive time stepping techniques which dynamically choose the step size needed for a given problem.

This research was conducted as part of an ASCI project. The principle investigators are David Keyes (ICASE and Old Dominion University), Dimitri Mavriplis (ICASE) and Alex Pothen (ICASE and Old Dominion University).

REPORTS AND ABSTRACTS

Girimaji, Sharath S.: *Reduction of large dynamical systems by minimization of evolution rate.* ICASE Report No. 99-15, (NASA/CR-1999-209121), May 19, 1999, 13 pages. To appear in Physical Review Letters.

Reduction of a large system of equations to a lower-dimensional system of similar dynamics is investigated. For dynamical systems with disparate timescales, a criterion for determining redundant dimensions and a general reduction method based on the minimization of evolution rate are proposed.

Janka, Erik, and Ponnampalam Balakumar: *On the stability of three-dimensional boundary layers part 1: Linear and nonlinear stability.* ICASE Report No. 99-16, (NASA/CR-1999-209330), May 28, 1999, 25 pages. Submitted to Theoretical and Computational Fluid Dynamics.

The primary stability of incompressible three-dimensional boundary layers is investigated using the Parabolized Stability Equations (PSE). We compute the evolution of stationary and traveling disturbances in the linear and nonlinear region prior to transition. As model problems, we choose Swept Hiemenz Flow and the DLR Transition Experiment. The primary stability results for Swept Hiemenz Flow agree very well with computations by Malik et al. For the DLR Experiment, the mean flow profiles are obtained by solving the boundary layer equations for the measured pressure distribution. Both linear and nonlinear results show very good agreement with the experiment.

Llorente, Ignacio M., Boris Diskin, and N. Duane Melson: *Plane smoothers for multiblock grids: Computational aspects.* ICASE Report No. 99-17, (NASA/CR-1999-209331), May 28, 1999, 29 pages. To be submitted to the Journal of Computational Physics.

Standard multigrid methods are not well suited for problems with anisotropic discrete operators, which can occur, for example, on grids that are stretched in order to resolve a boundary layer. One of the most efficient approaches to yield robust methods is the combination of standard coarsening with alternating-direction plane relaxation in the three dimensions. However, this approach may be difficult to implement in codes with multiblock structured grids because there may be no natural definition of global lines or planes. This inherent obstacle limits the range of an implicit smoother to only the portion of the computational domain in the current block. This report studies in detail, both numerically and analytically, the behavior of blockwise plane smoothers in order to provide guidance to engineers who use block-structured grids. The results obtained so far show alternating-direction plane smoothers to be very robust, even on multiblock grids. In common computational fluid dynamics multiblock simulations, where the number of subdomains crossed by the line of a strong anisotropy is low (up to four), textbook multigrid convergence rates can be obtained with a small overlap of cells between neighboring blocks.

Lüttgen, Gerald, and Victor Carreño: *Analyzing mode confusion via model checking.* ICASE Report No. 99-18, (NASA/CR-1999-209332), May 27, 1999, 67 pages. Submitted to the Sixth Spin99 Workshop.

Mode confusion is one of the most serious problems in aviation safety. Today's complex digital flight decks make it difficult for pilots to maintain awareness of the actual states, or modes, of the flight deck automation. NASA Langley leads an initiative to explore how formal techniques can be used to discover possible sources of mode confusion. As part of this initiative, a flight guidance system was previously specified as a finite Mealy automaton, and the theorem prover PVS was used to reason about it.

The objective of the present paper is to investigate whether state-exploration techniques, especially model checking, are better able to achieve this task than theorem proving and also to compare several verification tools for the specific application. The flight guidance system is modeled and analyzed in Murphi, SMV, and Spin. The tools are compared regarding their system description language, their practicality for analyzing mode confusion, and their capabilities for error tracing and for animating diagnostic information. It turns out that their strengths are complementary.

Janke, Erik, and Ponnampalam Balakumar: *On the stability of three-dimensional boundary layers part 2: Secondary instability.* ICASE Report No. 99-19, (NASA/CR-1999-209341), June 15, 1999, 28 pages. Submitted to Theoretical and Computational Fluid Dynamics.

The secondary instability of three-dimensional incompressible boundary layers is studied using Floquet theory. Starting from the equilibrium solutions that we obtained from the PSE computations documented in Part 1, we investigate the region where a purely stationary crossflow disturbance saturates for its secondary instability characteristics utilizing developed global and local eigenvalue solvers that are based on the Implicitly Restarted Arnoldi Method, and a Newton-Raphson technique, respectively. The main focuses of this study are on the existence of multiple roots in the eigenvalue spectrum that could explain experimental observations of time-dependent occurrences of an explosive growth of traveling disturbances, on the routes by which high-frequency disturbances enter the boundary layer, as well as on gaining more information about threshold amplitudes for the growth of secondary disturbances.

Sidilkover, David: *Factorizable schemes for the equations of fluid flows.* ICASE Report No. 99-20, (NASA/CR-1999-209345), June 15, 1999, 18 pages. Submitted to Applied Numerical Mathematics.

We present an upwind high-resolution factorizable (UHF) discrete scheme for the compressible Euler equations that allows to distinguish between full-potential and advection factors at the discrete level. The scheme approximates equations in their general conservative form and is related to the family of genuinely multidimensional upwind schemes developed previously and demonstrated to have good shock-capturing capabilities. A unique property of this scheme is that in addition to the aforementioned features it is also factorizable, i.e., it allows to distinguish between full-potential and advection factors at the discrete level. The latter property facilitates the construction of optimally efficient multigrid solvers. This is done through a relaxation procedure that utilizes the factorizability property.

Qian, Yue-Hong, and Ye Zhou: *On higher order dynamics in lattice-based models using Chapman-Enskog method.* ICASE Report No. 99-21, (NASA/CR-1999-209346), June 18, 1999, 12 pages. Submitted to Physical Review E.

In this paper, we investigate the existence of higher order dynamics in lattice-based models. We have identified two conditions that determine whether a model would allow some Burnett-like equations when the Chapman-Enskog expansion is used. These two conditions are the number of the conserved quantities as well as the space and time discretization. We shall demonstrate these conditions by discussing (1) pure diffusion equation, and (2) hydrodynamic equations. While the fact that diffusion equation allows the higher order dynamics can be shown easily, we will illustrate that care must be taken when deriving Burnett-like equations for lattice-based hydrodynamics models using the Chapman-Enskog method.

Kennedy, Christopher A., Mark H. Carpenter, and R. Michael Lewis: *Low-storage, explicit Runge-Kutta schemes for the compressible Navier-Stokes equations.* ICASE Report No. 99-22, (NASA/CR-1999-209349), June 18, 1999, 57 pages. Submitted to Applied Numerical Mathematics.

The derivation of low-storage, explicit Runge-Kutta (ERK) schemes has been performed in the context of integrating the compressible Navier-Stokes equations via direct numerical simulation. Optimization of ERK methods is done across the broad range of properties, such as stability and accuracy efficiency, linear and nonlinear stability, error control reliability, step change stability, and dissipation/dispersion accuracy, subject to varying degrees of memory economization. Following van der Houwen and Wray, 16 ERK pairs are presented using from two to five registers of memory per equation, per grid point and having accuracies from third- to fifth-order. Methods have been assessed using the differential equation testing code DETEST, and with the 1D wave equation. Two of the methods have been applied to the DNS of a compressible jet as well as methane-air and hydrogen-air flames. Derived 3(2) and 4(3) pairs are competitive with existing full-storage methods. Although a substantial efficiency penalty accompanies use of two- and three-register, fifth-order methods, the best contemporary full-storage methods can be nearly matched while still saving two to three registers of memory.

Ryaben'kii, V.S., V.I. Turchaninov, and S.V. Tsynkov: *Long-time numerical integration of the three-dimensional wave equation in the vicinity of a moving source.* ICASE Report No. 99-23, (NASA/CR-1999-209350), June 18, 1999, 25 pages. To be submitted to Mathematics of Computation.

We propose a family of algorithms for solving numerically a Cauchy problem for the three-dimensional wave equation. The sources that drive the equation (i.e., the right-hand side) are compactly supported in space for any given time; they, however, may actually move in space with a subsonic speed. The solution is calculated inside a finite domain (e.g., sphere) that also moves with a subsonic speed and always contains the support of the right-hand side.

The algorithms employ a standard consistent and stable explicit finite-difference scheme for the wave equation. They allow one to calculate the solution for arbitrarily long time intervals without error accumulation and with the fixed non-growing amount of the CPU time and memory required for advancing one time step. The algorithms are inherently three-dimensional; they rely on the presence of lacunae in the solutions of the wave equation in oddly dimensional spaces.

The methodology presented in the paper is, in fact, a building block for constructing the nonlocal highly accurate unsteady artificial boundary conditions to be used for the numerical simulation of waves propagating with finite speed over unbounded domains.

Alexandrov, Natalia M., and Robert Michael Lewis: *Comparative properties of collaborative optimization and other approaches to MDO.* ICASE Report No. 99-24, (NASA/CR-1999-209354), July 13, 1999, 17 pages. To appear in the Proceedings of ASMO UK/ISSMO Conference on Engineering Design Optimization.

We discuss criteria by which one can classify, analyze, and evaluate approaches to solving multidisciplinary design optimization (MDO) problems. Central to our discussion is the often overlooked distinction between questions of formulating MDO problems and solving the resulting computational problem. We illustrate our general remarks by comparing several approaches to MDO that have been proposed.

Diskin, Boris: *Solving upwind-biased discretizations II: Multigrid solver using semicoarsening.* ICASE Report No. 99-25, (NASA/CR-1999-209355), July 14, 1999, 34 pages. To be submitted to the SIAM Journal of Scientific Computing.

This paper studies a novel multigrid approach to the solution for a second order upwind biased discretization of the convection equation in two dimensions. This approach is based on semicoarsening and well balanced explicit correction terms added to coarse-grid operators to maintain on coarse grids the same cross-characteristic interaction as on the target (fine) grid. Colored relaxation schemes are used on all the levels allowing a very efficient parallel implementation. The results of the numerical tests can be summarized as follows:

1. The residual asymptotic convergence rate of the proposed $V(0,2)$ multigrid cycle is about 3 per cycle. This convergence rate far surpasses the theoretical limit (4/3) predicted for standard multigrid algorithms using full coarsening. The reported efficiency does not deteriorate with increasing the cycle depth (number of levels) and/or refining the target-grid mesh spacing.
2. The full multigrid algorithm (FMG) with two $V(0,2)$ cycles on the target grid and just one $V(0,2)$ cycle on all the coarse grids always provides an approximate solution with the algebraic error less than the discretization error. Estimates of the total work in the FMG algorithm are ranged between 18 and 30 minimal work units (depending on the target discretization). Thus, the overall efficiency of the FMG solver closely approaches (if does not achieve) the goal of the textbook multigrid efficiency.
3. A novel approach to deriving a discrete solution approximating the true continuous solution with a relative accuracy given in advance is developed. An adaptive multigrid algorithm (AMA) using comparison of the solutions on two successive target grids to estimate the accuracy of the current target-grid solution is defined. A desired relative accuracy is accepted as an input parameter. The final target grid on which this accuracy can be achieved is chosen automatically in the solution process. The actual relative accuracy of the discrete solution approximation obtained by AMA is always better than the required accuracy; the computational complexity of the AMA algorithm is (nearly) optimal (comparable with the complexity of the FMG algorithm applied to solve the problem on the optimally spaced target grid).

Muñoz, César, and John Rushby: *Structural embeddings: Mechanization with method.* ICASE Report No. 99-26, (NASA/CR-1999-209360), July 16, 1999, 22 pages. To appear in the Proceedings of FM'99.

The most powerful tools for analysis of formal specifications are general-purpose theorem provers and model checkers, but these tools provide scant methodological support. Conversely, those approaches that do provide a well-developed method generally have less powerful automation. It is natural, therefore, to try to combine the better-developed methods with the more powerful general-purpose tools. An obstacle is that the methods and the tools often employ very different logics.

We argue that methods are separable from their logics and are largely concerned with the structure and organization of specifications. We propose a technique called structural embedding that allows the structural elements of a method to be supported by a general-purpose tool, while substituting the logic of the tool for that of the method. We have found this technique quite effective and we provide some examples of its application. We also suggest how general-purpose systems could be restructured to support this activity better.

Liu, Jian-Guo, and Chi-Wang Shu: *A high order discontinuous Galerkin method for 2D incompressible flows.* ICASE Report No. 99-27, (NASA/CR-1999-209361), July 16, 1999, 22 pages. Submitted to the Journal of Computational Physics.

In this paper we introduce a high order discontinuous Galerkin method for two dimensional incompressible flow in vorticity streamfunction formulation. The momentum equation is treated explicitly, utilizing the efficiency of the discontinuous Galerkin method. The streamfunction is obtained by a standard Poisson solver using continuous finite elements. There is a natural matching between these two finite element spaces, since the normal component of the velocity field is continuous across element boundaries. This allows for a correct upwinding gluing in the discontinuous Galerkin framework, while still maintaining total energy conservation with no numerical dissipation and total enstrophy stability. The method is suitable for inviscid or high Reynolds number flows. Optimal error estimates are proven and verified by numerical experiments.

Lian, Yongsheng, and Kun Xu: *Gas-kinetic scheme for multimaterial flows and its application in chemical reaction.* ICASE Report No. 99-28, (NASA/CR-1999-209364), August 20, 1999, 34 pages. To be submitted to the Journal of Computational Physics.

This paper concerns the extension of the multicomponent gas-kinetic BGK-type scheme to multidimensional chemical reactive flow calculations. In the kinetic model, each component satisfies its individual gas-kinetic BGK equation and the equilibrium states of both components are coupled in space and time due to the momentum and energy exchange in the course of particle collisions. At the same time, according to the chemical reaction rule one component can be changed into another component with the release of energy, where the reactant and product could have different gamma. Many numerical test cases are included in this paper, which show the robustness and accuracy of kinetic approach in the description of multicomponent reactive flows.

Smith, Ralph C., and Zoubeida Ounaies: *A hysteresis model for piezoceramic materials.* ICASE Report No. 99-29, (NASA/CR-1999-209368), August 20, 1999, 13 pages. To be presented at the 1999 ASME International Mechanical Engineering Congress and Exposition.

This paper addresses the modeling of nonlinear constitutive relations and hysteresis inherent to piezoceramic materials at moderate to high drive levels. Such models are necessary to realize the full potential of the materials in high performance control applications, and a necessary prerequisite is the development of techniques which permit control implementation. The approach employed here is based on the quantification of reversible and irreversible domain wall motion in response to applied electric fields. A comparison with experimental data illustrates that because the resulting ODE model is physic-based, it can be employed for both characterization and prediction of polarization levels throughout the range of actuator operation. Finally, the ODE formulation is amenable to inversion which facilitates the development of an inverse compensator for linear control design.

Yefet, Amir, and Peter G. Petropoulos: *A non-dissipative staggered fourth-order accurate explicit finite difference scheme for the time-domain Maxwell's equations.* ICASE Report No. 99-30, (NASA/CR-1999-209514), August 20, 1999, 30 pages. To be submitted to the SIAM Journal of Applied Mathematics.

We consider a divergence-free non-dissipative fourth-order explicit staggered finite difference scheme for the hyperbolic Maxwell's equations. Special one-sided difference operators are derived in order to implement the scheme near metal boundaries and dielectric interfaces. Numerical results show the scheme is long-time stable, and is fourth-order convergent over complex domains that include dielectric interfaces and perfectly conducting surfaces. We also examine the scheme's behavior near metal surfaces that are not alligned with the grid axes, and compare its accuracy to that obtained by the Yee scheme.

Xu, Kun: *A gas-kinetic method for the hyperbolic-elliptic equations and its application in two-phase fluid flow.* ICASE Report No. 99-31, (NASA/CR-1999-209515), August 20, 1999, 20 pages. To be submitted to the SIAM Journal of Scientific Computing.

A gas-kinetic method for the hyperbolic-elliptic equations is presented in this paper. In the mixed type system, the co-existence and the phase transition between liquid and gas are described by the van der Waals-type equation of state (EOS). Due to the unstable mechanism for a fluid in the elliptic region, interface between the liquid and gas can be kept sharp through the condensation and evaporation process to remove the "averaged" numerical fluid away from the elliptic region, and the interface thickness depends on the numerical diffusion and stiffness of the phase change. A few examples are presented in this paper for both phase transition and multifluid interface problems.

Ounaies, Zoubeida, Cheol Park, Joycelyn S. Harrison, Joseph G. Smith, and Jeffrey Hinkley: *Structure-property study of piezo-electricity in polyimides.* ICASE Report No. 99-32, (NASA/CR-1999-209516), August 20, 1999, 14 pages. To appear in the Proceedings for the SPIE: Smart Structures and Materials Symposium.

High performance piezoelectric polymers are of interest to NASA as they may be useful for a variety of sensor applications. Over the past few years research on piezoelectric polymers has led to the development of promising high temperature piezoelectric responses in some novel polyimides. In this study, a series of polyimides have been studied with systematic variations in the diamine monomers that comprise the polyimide while holding the dianhydride constant. The effect of structural changes, including variations in the nature and concentration of dipolar groups, on the remanent polarization and piezoelectric coefficient is examined. Fundamental structure-piezoelectric property insight will enable the molecular design of polymers possessing distinct improvements over state-of-the-art piezoelectric polymers including enhanced polarization, polarization stability at elevated temperatures, and improved processability.

Demuren, Ayodeji O., and Robert V. Wilson: *Streamwise vorticity generation in laminar and turbulent jets.* ICASE Report No. 99-33, (NASA/CR-1999-209517), August 20, 1999, 16 pages. To appear in the Proceedings of ASME/JSME Fluids Engineering Conference.

Complex streamwise vorticity fields are observed in the evolution of non-circular jets. Generation mechanisms are investigated via Reynolds-averaged (RANS), large-eddy (LES) and direct numerical (DNS) simulations of laminar and turbulent rectangular jets. Complex vortex interactions are found in DNS of laminar jets, but axis-switching is observed only when a single instability mode is present in the incoming mixing layer. With several modes present, the structures are not coherent and no axis-switching occurs. RANS computations also produce no axis-switching. On the other hand, LES of high Reynolds number turbulent jets produce axis-switching even for cases with several instability modes in the mixing layer. Analysis of the source terms of the mean streamwise vorticity equation through post-processing of the instantaneous results shows that complex interactions of gradients of the normal and shear Reynolds stresses are responsible for the generation of streamwise vorticity which leads to axis-switching. RANS computations confirm these results. k-ϵ turbulence model computations fail to reproduce the phenomenon, whereas algebraic Reynolds stress model (ASM) computations, in which the secondary normal and shear stresses are computed explicitly, succeeded in reproducing the phenomenon accurately.

Povitsky, Alex, and Philip J. Morris: *A parallel compact multi-dimensional numerical algorithm with aeroacoustics applications.* ICASE Report No. 99-34, (NASA/CR-1999-209518), August 23, 1999, 24 pages. Presented at the 14th AIAA CFD Conference.

In this study we propose a novel method to parallelize high-order compact numerical algorithms for the solution of three-dimensional PDEs in a space-time domain. For this numerical integration most of the computer time is spent in computation of spatial derivatives at each stage of the Runge-Kutta temporal update. The most efficient direct method to compute spatial derivatives on a serial computer is a version of Gaussian elimination for narrow linear banded systems known as the Thomas algorithm. In a straightforward pipelined implementation of the Thomas algorithm processors are idle due to the forward and backward recurrences of the Thomas algorithm. To utilize processors during this time, we propose to use them for either non-local data independent computations, solving lines in the next spatial direction, or local data-dependent computations by the Runge-Kutta method. To achieve this goal, control of processor communication and computations by a static schedule is adopted. Thus, our parallel code is driven by a communication and computation schedule instead of the usual "creative programming" approach. The obtained parallelization speed-up of the novel algorithm is about twice as much as that for the standard pipelined algorithm and close to that for the explicit DRP algorithm.

Baggag, Abdelkader, Harold Atkins, and David Keyes: *Parallel implementation of the discontinuous Galerkin method.* ICASE Report No. 99-35, (NASA/CR-1999-209546), August 31, 1999, 12 pages. Submitted to the Proceedings of the International Parallel CFD 1999 Conference.

This paper describes a parallel implementation of the discontinuous Galerkin method. Discontinuous Galerkin is a spatially compact method that retains its accuracy and robustness on non-smooth unstructured grids and is well suited for time dependent simulations. Several parallelization approaches are studied and evaluated. The most natural and symmetric of the approaches has been implemented in an object-oriented code used to simulate aeroacoustic scattering. The parallel implementation is MPI-based and has been tested on various parallel platforms such as the SGI Origin, IBM SP2, and clusters of SGI and Sun workstations. The scalability results presented for the SGI Origin show slightly superlinear speedup on a fixed-size problem due to cache effects.

Roe, Kevin, and Piyush Mehrotra: *Parallelization of a multigrid incompressible viscous cavity flow solver using OpenMP.* ICASE Report No. 99-36, (NASA/CR-1999-209551), September 13, 1999, 13 pages. Proceedings of the Ninth SIAM Conference on Parallel Processing for Scientific Computing.

We describe a multigrid scheme for solving the viscous incompressible driven cavity problem that has been parallelized using OpenMP. The incremental parallelization allowed by OpenMP was of great help during the parallelization process. Results show good parallel efficiencies for reasonable problem sizes on an SGI Origin 2000. Since OpenMP allowed us to specify the number of threads (and in turn processors) at runtime, we were able to improve performance when solving on smaller/coarser meshes. This was accomplished by giving each processor a more reasonable amount of work rather than having many processors work on very small segments of the data (and thereby adding significant overhead).

Lasseigne, D.G., W.O. Criminale, R.D. Joslin, and T.L. Jackson: *Towards understanding the mechanism of receptivity and bypass dynamics in laminar boundary layers.* ICASE Report No. 99-37, (NASA/CR-1999-209553), September 13, 1999, 31 pages. To be submitted to the Journal of Fluid Mechanics.

Three problems concerning Laminar-Turbulent Transition are addressed by solving a series of initial value problems. The first problem is the calculation of resonance within the continuous spectrum of the Blasius boundary layer. The second is calculation of the growth of Tollmien-Schlichting waves that are a direct result of disturbances that only lie outside of the boundary layer. And, the third problem is the calculation of non-parallel effects. Together, these problems represent a unified approach to the study of freestream disturbance effects that could lead to transition. Solutions to the temporal, initial-value problem with an inhomogeneous forcing term imposed upon the flow is sought. By solving a series of problems, it is shown that:

- A transient disturbance lying completely outside of the boundary layer can lead to the growth of an unstable Tollmien-Schlichting wave.
- A resonance with the continuous spectrum leads to strong amplification that may provide a mechanism for bypass transition once nonlinear effects are considered.
- A disturbance with a very weak unstable Tollmien-Schlichting wave can lead to a much stronger Tollmien-Schlichting wave downstream, if the original disturbance has a significant portion of its energy in the continuum modes.

Dobrian, Florin, Gary Kumfert, and Alex Pothen: *The design of sparse direct solvers using object-oriented techniques.* ICASE Report No. 99-38, (NASA/CR-1999-209558), September 15, 1999, 37 pages. To appear in Modern Software Tools in Scientific Computing.

We describe our experience in designing object-oriented software for sparse direct solvers. We discuss Spindle, a library of sparse matrix ordering codes, and OBLIO, a package that implements the factorization and triangular solution steps of a direct solver. We discuss the goals of our design: managing complexity, simplicity of interface, flexibility, extensibility, safety, and efficiency. High performance is obtained by carefully implementing the computationally intensive kernels and by making several tradeoffs to balance the conflicting demands of efficiency and good software design. Some of the missteps that we made in the course of this work are also described.

INTERIM REPORTS

Chen, Po-Shu: *Implementation of interaction algorithm to non-matching discrete interfaces between structure and fluid mesh.* ICASE Interim Report No. 36, (NASA/CR-1999-209340), June 18, 1999, 13 pages.

This paper presents software for solving the non-conforming fluid structure interfaces in aeroelastic simulation. It reviews the algorithm of interpolation and integration, highlights the flexibility and the user-friendly feature that allows the user to select the existing structure and fluid package, like NASTRAN and CLF3D, to perform the simulation. The presented software is validated by computing the High Speed Civil Transport model.

ICASE COLLOQUIA

April 1, 1999 – September 30, 1999

Name/Affiliation/Title	Date
Shaffer, Cliff, Virginia Polytechnic Institute and State University "Visualization for Multiparameter Aircraft Designs"	April 5
Filippova, Olga, Gerhard-Mercator-Universitaet, Germany "Lattice-BGK Model for Low Mach Number Flows"	April 6
Haftka, Raphael, University of Florida, Gainesville "Development of Inexpensive Experiments for Testing Methods for Design Against Uncertainty"	April 9
Huyse, Luc, University of Calgary, Canada "The Imperative of Including Uncertainties in a Design"	April 12
Rossow, Cord-Christian, German Aerospace Center "The German National CFD Initiative MEGAFLOW"	April 20
Kogan, Mikhail, TsAGI - Central Aerohydrodynamics Institute, Russia "TsAGI Research on the Influence of Aviation on the Atmosphere"	April 28
Dubois, Evelyne, Nice University, France "A Novel Sensitivity Analysis Method for High Fidelity Multidisciplinary Optimization of Aero-structural Systems"	May 3
Fahl, Marco, University of Trier, Germany "On the Application of POD in Flow Control"	May 6
Hesthaven, Jan, Brown University "Spectral Methods on Unstructured Grids - Accurate, Stable, and Fast"	May 7
Ciardo, Gianfranco, The College of William & Mary "Efficient State Space Exploration and Storage Using Decision Diagrams"	May 10
Mendler, Michael, Sheffield University "Lax Logic and Its Application to the Timing Analysis of Combinational Systems"	May 18

Name/Affiliation/Title	Date
Giunta, Anthony A., National Research Council, NASA Langley "A Novel Sensitivity Analysis Method for High Fidelity Multidisciplinary Optimization of Aero-structural Systems"	May 21
Brenner, Gunther, Institute of Fluid Mechanics, Erlangen University, Germany "Modeling and Simulation of CFD Problems Related to Process and Chemical Engineering"	May 27
Naguib, Ahmed, Michigan State "An Investigation of Wall-pressure Flow Sources in a Turbulent Boundary Layer"	June 18
Epstein, Boris, The Academic College of Tel Aviv - Yaffo "Aerodynamically Accurate Navier-Stokes Computations"	June 22
Cleaveland, Rance, State University of New York at Stony Brook "Verifying Active Structural Control Systems: A Case Study in Formal Analysis"	June 29
Shusser, Michael, California Institute of Technology "Inviscid Vortex Ring Formation"	July 2
Roska, Tamas, Computer and Automation Institute Budapest, Hungary "Analogic Cellular Computing and a TeraOPS Speed Focal Plane Visual Microprocessor"	July 13
Schneider, Frank, California Institute of Technology "Model Checking in the Context of the Spin Validation System"	July 15
Brown, Garry L., Princeton University "Experiments on Stability, Receptivity and Transition in a Mach 3 Boundary Layer"	July 21
Goossens, Serge, Katholieke Universitaeit, Belgium "Two-level Algorithms for Overlapping Composite Mesh Difference Methods"	July 22
Taft, James R., Silicon Graphics, NASA Ames "Recent CFD Performance Results Using the New 02K with R12000 Processor"	July 23
Kopetz, Hermann, Technische Universitaet Wien, Austria "The Time-triggered Model of Computation"	July 26
Hyland, David, University of Michigan, Ann Arbor "Toward Self-reliant Control"	July 27

Name/Affiliation/Title	Date
Turner, Leaf, Los Alamos National Laboratory "Helicity Decomposition: A Compact Descriptor for Inhomogeneous Turbulence Dynamics"	July 29
Povitsky, Alex, ICASE "High-order Compact Numerical Schemes and Modern Computers"	July 30
Venditti, David, Massachusetts Institute of Technology "Adjoint Error Estimation and Grid Adaptive Criteria for Accurate CFD Predictions of Integral Outputs"	July 30
Habermehl, Peter, University of Paris "Verification of Infinite-state Systems by Combining Abstraction and Reachability Analysis"	August 3
Mistree, Farrokh, Georgia Institute of Technology "Product Realization in a Distributed Engineering Environment"	August 11
Yefet, Amir, New Jersey Institute of Technology "A Non-dissipative Staggered Fourth-order Accurate Explicit Finite Difference Scheme for the Time-domain Maxwell's Equations"	August 11
Rossow, Cord-Christian, German Aerospace Center "On Implicit Residual Smoothing Techniques"	August 17
Johnson, Chris, University of Glasgow "Next Generation Accident Reporting"	August 25
Kulkarni, Sandeep, Ohio State University "Component Based Design of Fault-tolerance"	August 27
Bokhari, Shahid, Pakistan University of Engineering "Experience with the Tera MTA"	August 30
Karpel, Moti, Technion - Israel Institute of Technology "Reduced-size Models for Integrated Aeroservoelastic Analysis and Design Optimization"	August 30
Raveh, Daniella, Georgia Institute of Technology "Reduced-sized Computational Aeroelastic Models for Aircraft Design Optimization"	August 31
Roe, Philip, University of Michigan, Ann Arbor "Finite Volume Schemes that Preserve Vorticity"	September 1

Name/Affiliation/Title	Date
Roe, Philip, University of Michigan, Ann Arbor "Euler Codes Giving Potential Flow and Boltzmann Codes Giving Euler Flow"	September 2
Aksay, Ilhan, Princeton University "Biomimetics Seminar: Bioinspired Processing of Organic/Inorganic Composites through Self-assembly"	September 3
Bjorner, Nikolaj, Kestrel Institute, Palo Alto, CA "Integrating Decision Procedures for Temporal Verification"	September 3
Dickinson, Michael, University of California, Berkeley "Biomimetics Seminar: The Aerodynamic Basis of Insect Flight"	September 13
Carman, Greg P., University of California, Los Angeles "Biomimetics Seminar: Active Materials Research at UCLA"	September 15
Hamba, Fujihiro, University of Tokyo "Effects of Pressure Fluctuations on Turbulence Growth in Compressible Homogeneous Shear Flow"	September 21
Rapoff, Andrew J., University of Florida, Gainesville "Biomimetics Seminar: Biomimetic Structures and Materials Research at the University of Florida"	September 29

ICASE SUMMER ACTIVITIES

The summer program for 1999 included the following visitors:

VISITOR and AREA OF RESEARCH	AFFILIATION	DATE OF VISIT
Abarbanel, Saul *Applied & Numerical Math*	Tel Aviv University, Israel	7/20 – 8/13
Al-Theneyan, Ahmed *Computer Science*	Old Dominion University	5/24 – 8/27
Baggag, Abdelkader *Computer Science*	Purdue University	5/24 – 9/10
Banks, H. Thomas *Applied & Numerical Math*	North Carolina State University	6/08 – 6/09 7/06 – 7/07 8/31 – 9/03
Bjorner, Nikolaj *Applied & Numerical Math*	Kestrel Institute	8/30 – 9/03
Brandt, Achi *Applied & Numerical Math*	The Weizmann Institute of Science, Israel	7/05 – 7/16
Chapman, Barbara *Computer Science*	University of Houston	8/09 – 8/20
Ciardo, Gianfranco *Computer Science*	The College of William & Mary	7/19 – 8/13
Criminale, William *Physical Sciences – Fluid Mechanics*	University of Washington, Seattle	9/20 – 9/24
Cronk, David *Computer Science*	Lucent Technologies	7/19 – 7/23
Darmofal, David *Applied & Numerical Math*	Massachusetts Institute of Technology	8/23 – 9/03

VISITOR and AREA OF RESEARCH	AFFILIATION	DATE OF VISIT
Fu, Jinghua *Computer Science*	Old Dominion University	6/01 - 10/7
Gottlieb, David *Applied & Numerical Math*	Brown University	6/28 - 7/01 7/26 - 7/29 8/23 - 8/25
Grosch, Chester *Physical Sciences - Fluid Mechanics*	Old Dominion University	5/03 - 5/07 7/05 - 7/30
Habermehl, Peter *Computer Science*	University of Paris 7	8/02 - 8/13
Huston, Kerry *Physical Sciences - Flow Control*	Virginia Polytechnic Institute and State University	5/24 - 8/13
Jackson, Tom *Physical Sciences - Fluid Mechanics*	University of Illinois, Urbana	9/20 - 9/24
Kapania, Rakesh *Applied & Numerical Math*	Virginia Polytechnic Institute and State University	8/09 - 8/20 9/20 - 9/24
Keyes, David *Computer Science*	Old Dominion University	5/17 - 5/28 7/12 - 7/23
Kopetz, Herman *Computer Science*	Technische Universitaet Wien, Austria	7/23 - 7/28
Lallemand, Pierre *Computer Science*	Centre National de la Recherche Scientifique, France	8/01 - 9/24
Lin, Chuang *Physical Sciences - Flow Control*	National Tsing Hua University	8/09 - 2/28/00
Llorente, Ignacio *Applied & Numerical Math*	Universidad Complutense, Spain	7/05 - 7/30
Longman, Richard *Physical Sciences - Flow Control*	Columbia University	6/07 - 7/09 8/02 - 9/03

VISITOR and AREA OF RESEARCH	AFFILIATION	DATE OF VISIT
Mahadevan, Sankaran *Applied & Numerical Math*	Vanderbilt University	8/01 – 8/31
Mattsson, Ken *Applied & Numerical Math*	Uppsala University, Sweden	5/03 5/28
Mendler, Michael *Computer Science*	Sheffield University, United Kingdom	5/14 – 6/03
Milder, Seth *Computer Science*	George Mason University	6/07 – 8/13
Montero, Ruben Santiago *Applied & Numerical Math*	Universidad Complutense, Spain	7/05 – 8/27
Naguib, Ahmed M. *Physical Sciences Flow Control*	Michigan State	6/07 – 7/02
Nash, Stephen *Applied & Numerical Math*	George Mason University	8/09 8/13
Nordstrom, Jan *Applied & Numerical Math*	The Aeronautical Research Institute of Sweden	6/21 7/09
Prieto, Manuel *Applied & Numerical Math*	Universidad Complutense de Madrid, Spain	7/12 – 8/13
Roe, Philip *Applied & Numerical Math*	University of Michigan, Ann Arbor	8/30 – 9/03
Schneider, Frank *Computer Science*	California Institute of Technology	7/12 7/16
Shu, Chi-Wang *Physical Sciences – Fluid Mechanics*	Brown University	8/16 – 8/30
Siminiceanu, Radu *Computer Science*	The College of William & Mary	7/19 8/13

VISITOR and AREA OF RESEARCH	AFFILIATION	DATE OF VISIT
Tsynkov, Semyon *Applied & Numerical Math*	Tel Aviv University, Israel	6/28 - 9/24
Turkel, Eli *Applied & Numerical Math*	Tel Aviv University, Israel	6/28 - 8/06
van Leer, Bram *Applied & Numerical Math*	University of Michigan, Ann Arbor	9/07 - 9/10
Venditti, David *Applied & Numerical Math*	Massachusetts Institute of Technology	7/02 - 7/30
Xu, Kun *Applied & Numerical Math*	The Hong Kong University of Science and Technology	6/10 - 8/13
Yefet, Amir *Applied & Numerical Math*	New Jersey Institute of Technology	8/02 - 8/13

OTHER ACTIVITIES

On April 27, 1999, ICASE and NASA Langley Research Center co-sponsored the Third Biennial Theodorsen Lectureship Award. Mark V. Morkovin, Professor Emeritus, Illinois Institute of Technology, was presented this award for his lifetime contributions to aerodynamics and turbulence research. His award lecture entitled "On Fluid Mechanics, Instabilities, Turbulence and Chaos" was presented at the Pearl Young Theater with over 100 in attendance.

On May 23–26, 1999, ICASE, NASA Langley Research Center, Old Dominion University, and The College of William and Mary co-sponsored the Parallel CFD'99 Conference at the Williamsburg Hospitality House in Williamsburg, VA. This conference featured many diverse realms of phenomena in which fluid dynamical simulations play a critical role. There were 125 attendees, and a formal proceedings will be published by Elsevier.

On June 22–25, 1999, CEAS, AIAA, ICASE, and NASA Langley Research Center co-sponsored the International Forum on Aeroelasticity and Structural Dynamics 1999, held at the Williamsburg Hospitality House in Williamsburg, VA. This Forum allowed scientists and engineers from industry, government, and universities to exchange knowledge and results of current studies and to discuss directions for future research and development for aircraft and spacecraft dynamics. There were 176 participants, and NASA Langley published the proceedings.

ICASE STAFF

I. ADMINISTRATIVE

Manuel D. Salas, Director, M.S., Aeronautics and Astronautics, Polytechnic Institute of Brooklyn, 1970. Fluid Mechanics and Numerical Analysis.

Linda T. Johnson, Office and Financial Administrator

Barbara A. Cardasis, Administrative Secretary

Etta M. Morgan, Accounting Supervisor

Emily N. Todd, Conference Manager/Executive Assistant

Shannon K. Verstynen, Information Technologist

Gwendolyn W. Wesson, Contract Accounting Clerk

Shouben Zhou, Systems Manager

Peter J. Kearney, Student Assistant

II. SCIENCE COUNCIL

David Gottlieb, (Chair) Professor, Division of Applied Mathematics, Brown University.

Lee Beach, Professor, Department of Physics, Computer Science & Engineering, Christopher Newport University.

Francine Berman, Professor, Department of Computer Science & Engineering, University of California-San Diego.

Joseph E. Flaherty, Amos Eaton Professor, Departments of Computer Science and Mathematical Sciences, Rensselaer Polytechnic Institute.

Geoffrey Fox, Director, Northeast Parallel Architectural Center, Syracuse University.

Forrester Johnson, Aerodynamics Research, Boeing Commercial Airplane Group.

Robert W. MacCormack, Professor, Department of Aeronautics and Astronautics, Stanford University.

Stanley G. Rubin, Professor, Department of Aerospace Engineering and Engineering Mechanics, University of Cincinnati.

Manuel D. Salas, Director, Institute for Computer Applications in Science and Engineering, NASA Langley Research Center.

III. RESEARCH FELLOWS

Dimitri Mavriplis - Ph.D., Mechanical and Aerospace Engineering, Princeton University, 1988. Applied & Numerical Mathematics [Grid Techniques for Computational Fluid Dynamics]. (February 1997 to August 2001)

Piyush Mehrotra - Ph.D., Computer Science, University of Virginia, 1982. Computer Science [Programming Languages for Multiprocessor Systems]. (January 1991 to September 1999)

IV. SENIOR STAFF SCIENTISTS

Thomas W. Crockett - B.S., Mathematics, The College of William & Mary, 1977. Computer Science [System Software for Parallel Computing, Computer Graphics, and Scientific Visualization]. (February 1987 to August 2000)

R. Michael Lewis - Ph.D., Mathematical Sciences, Rice University, 1989. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (May 1995 to August 2000)

Josip Lončarić - Ph.D., Applied Mathematics, Harvard University, 1985. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (March 1996 to August 2001)

Kwan-Liu Ma - Ph.D., Computer Science, University of Utah, 1993. Computer Science [Visualization]. (May 1993 to September 1999)

Robert Rubinstein - Ph.D., Mathematics, Massachusetts Institute of Technology, 1972. Fluid Mechanics [Turbulence Modeling, Turbulence Management, and Acoustics]. (May 1998 to August 2001)

David Sidilkover - Ph.D., Applied Mathematics, The Weizmann Institute of Science, 1989. Applied & Numerical Mathematics [Numerical Analysis and Algorithms]. (November 1994 to November 1999)

V. SCIENTIFIC STAFF

Brian G. Allan - Ph.D., Mechanical Engineering, University of California at Berkeley, 1996. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (February 1996 to November 1999)

Eyal Arian - Ph.D., Applied Mathematics, The Weizmann Institute of Science, Israel, 1995. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (October 1994 to August 1999)

Po-Shu Chen - Ph.D., Aerospace Engineering, University of Colorado-Boulder, 1997. Physical Sciences [Computational Structures]. (January 1998 to October 1999)

Roger C. Hart - Ph.D., Physics, University of Tennessee, 1991. Physical Sciences [Measurement Science and Technology]. (December 1998 to October 1999)

Gerald Lüttgen - Ph.D., Computer Science, University of Passau, Germany, 1998. Computer Science [Formal Methods Research for Safety Critical Systems]. (October 1998 to August 2000)

Li-Shi Luo - Ph.D., Physics, Georgia Institute of Technology, 1993. Computer Science [Parallel Algorithms]. (November 1996 to October 1999)

Cesar A. Muñoz - Ph.D., Computer Science, University of Paris 7, 1997. Computer Science [Formal Methods Research for Safety Critical Systems]. (May 1999 to April 2001)

Zoubeida Ounaies - Ph.D., Engineering Science and Mechanics, The Pennsylvania State University, 1995. Physical Sciences [Characterization of Advanced Piezoelectric Materials]. (March 1999 to November 1999)

Alexander Povitsky - Ph.D., Mechanical Engineering, Moscow Institute of Steel and Alloys Technology (MISA), Russia, 1988. Computer Science [Parallelization and Formulation of Higher Order Schemes for Aeroacoustics Noise Propagation]. (October 1997 to August 2000)

VI. VISITING SCIENTISTS

Sang-Hyon Chu - Ph.D., Chemical Engineering, Seoul National University, 1998. Physical Sciences [Smart Materials and Flow Control]. (March 1998 to July 1999)

Boris Diskin - Ph.D., Applied Mathematics, The Weizmann Institute of Science, Israel, 1998. Teaching Assistant, The Weizmann Institute of Science, Israel. Applied & Numerical Mathematics [Convergence Acceleration]. (July 1998 to September 2000)

Chaung Lin - Ph.D., Nuclear Engineering, University of California-Berkeley, 1983. Professor, Department of Engineering and System Science, National Tsing Hua University, Taiwan. Physical Sciences [Flow Modeling from Experimental Data and Real-time Control Using Predictive Control Techniques]. (August 1999 to February 2000)

Cord-Christian Rossow - Ph.D., Aerospace Engineering, Technical University of Braunschweig, Germany, 1988. Branch Head, Dr.-Ing, DLR, Institute of Design Aerodynamics, Germany. Applied & Numerical Mathematics. (February 1999 to August 1999)

David R. Picasso - M.S., Management, Stanford University, 1993. NASA Retired. Computer Science. (January 1999 to July 1999)

Linda Stals - Ph.D., Mathematics, Australian National University, 1996. Post-Doc, Department of Computer Science, Old Dominion University. Computer Science [Parallel Implicit Multilevel Algorithms]. (November 1998 to October 1999)

Nail K. Yamaleev - Ph.D., Numerical Methods and Mathematical Modeling, Moscow Institute of Physics and Technology, 1993. Senior Research Scientist, Department of Computational Mathematics, Institute of Mathematics, Ufa, Russia. Physical Science [Fluid Mechanics]. (February 1999 to September 1999)

Ye Zhou - Ph.D., Physics, The College of William & Mary, 1987. Department of Aerospace Science Engineering, Tuskegee University. Fluid Mechanics [Lattice Boltzmann Method to Gas Liquid Flow]. (October 1998 to May 1999)

VII. SHORT-TERM VISITING SCIENTISTS

Dov S. Bai - Ph.D., Applied Mathematics, Weizmann Institute of Science, 1985. Research Scientist, Wright-Patterson Air Force Base. Applied & Numerical Mathematics. (July 1999)

Nikolaj S. Bjorner - Ph.D., Computer Science, Stanford University, 1998. Computer Scientist, Kestrel Institute. Computer Science. (August 1999 to September 1999)

Barbara Chapman - M.S., Mathematics, University of Canterbury, Christchurch, New Zealand, 1985. Assistant Professor, Department of Computer Science, University of Houston. Computer Science. (August 1999)

Gianfranco Ciardo - Ph.D., Computer Science, Duke University, 1989. Assistant Professor, The College of William & Mary. Computer Science [Formal Methods]. (July 1999 to August 1999)

William O. Criminale - Ph.D., Aeronautics, The Johns Hopkins University, 1960. Professor, Department of Applied Mathematics, University of Washington. Fluid Mechanics. (September 1999)

David L. Darmofal - Ph.D., Aerospace Engineering, Massachusetts Institute of Technology, 1993. Assistant Professor, Department of Aeronautics and Astronautics, Massachusetts Institute of Technology. Applied & Numerical Mathematics [Convergence Acceleration]. August 1999 to September 1999.

Peter Habermehl - Ph.D., Computer Science, VERIMAG, University Joseph Fourier, Grenoble, France, 1998. Assistant Professor, LIAFA, University of Paris 7, France. Computer Science. (August 1999)

Thomas L. Jackson - Ph.D., Mathematics, Rensselaer Polytechnic Institute, 1985. Senior Research Scientist, Center for Simulation of Advanced Rockets, University of Illinois. Fluid Mechanics [Laminar Flow Control]. (September 1999)

Rakesh K. Kapania - Ph.D., Aerospace, Purdue University, 1985. Professor, Department of Aerospace & Ocean Engineering, Virginia Polytechnic Institute and State University. Applied & Numerical Mathematics [Multidisciplinary Optimization]. (August 1999)

Herman Kopetz - Ph.D., Physics, University of Vienna, Austria, 1968. Professor, Department of Computer Science, Technical University of Vienna, Austria. Computer Science [Formal Methods]. (July 1999)

Pierre Lallemand - Ph.D., Physics, Universite de Paris, 1966. Director of Research, Centre National de la Recherche Scientifique, A.S.C.I., Universite Paris-Sud. Computer Science. (August 1999 to September 1999)

Ignacio M. Llorente - Ph.D., Computer Science, Complutense University of Madrid, Spain, 1995. Associate Professor, Department of Computer Architecture, Complutense University of Madrid, Spain. Applied & Numerical Mathematics [Convergence Acceleration]. (July 1999)

Sankaran Mahadevan - Ph.D., Civil Engineering, Georgia Institute of Technology, 1988. Associate Professor, Department of Civil Engineering, Vanderbilt University. Applied & Numerical Mathematics [Managing Uncertainties]. (August 1999)

Michael Mendler - Ph.D., Computer Science, University of Edinburgh, 1993. Associate Professor, Department of Computer Science, University of Sheffield, United Kingdom. Computer Science [Formal Methods]. (May 1999)

Ahmed M. Naguib - Ph.D., Mechanical and Aerospace Engineering, Illinois Institute of Technology, 1992. Assistant Professor, Department of Mechanical Engineering, Michigan State University. Physical Sciences [Surface-mounted Sensor for Active Flow Control]. (June 1999 to July 1999)

Stephen G. Nash - Ph.D., Computer Science, Stanford University, 1982. Associate Dean, School of Information Technology and Engineering, George Mason University. Applied & Numerical Mathematics [Managing Uncertainties]. (August 1999)

Philip Roe - Ph.D., Aeronautics, University of Cambridge, United Kingdom, 1962. Professor, Department of Aerospace Engineering, University of Michigan. Applied & Numerical Mathematics. (August 1999 to September 1999)

Francis L. Schneider - Ph.D., Elementary Particle Physics, Ohio University, 1974. Member of Engineering Staff Senior Multidisciplinary, Jet Propulsion Laboratory, California Institute of Technology. Computer Science [Formal Methods]. (July 1999)

Eli Turkel - Ph.D., Applied Mathematics, New York University, 1970. Associate Professor, Department of Applied Mathematics, Tel Aviv University, Israel. Applied & Numerical Mathematics [Convergence Acceleration]. (June 1999 to August 1999)

Bram van Leer - Ph.D., Theoretical Astrophysics, Leiden State University, The Netherlands, 1970. Professor, Department of Aerospace Engineering, University of Michigan. Applied & Numerical Mathematics [Convergence Acceleration]. (September 1999)

Kun Xu - Ph.D., Astrophysics, Columbia University, 1993. Assistant Professor, Department of Mathematics, The Hong Kong University of Science and Technology, Hong Kong. Applied & Numerical Mathematics [Developing Gas Kinetic Schemes]. (June 1999 to August 1999)

Amir Yefet - Ph.D., Applied Mathematics, Tel Aviv University, 1999. Post-doc, Department of Mathematical Sciences, new Jersey Institute of Technology. Applied & Numerical Mathematics [Computational Electromagnetics]. (August 1999)

VIII. ASSOCIATE RESEARCH FELLOW

David E. Keyes - Ph.D., Applied Mathematics, Harvard University, 1984. Computer Science [Parallel Numerical Algorithms]

IX. CONSULTANTS

Saul Abarbanel - Ph.D., Theoretical Aerodynamics, Massachusetts Institute of Technology, 1959. Professor, Department of Applied Mathematics, Tel Aviv University, Israel. Applied & Numerical Mathematics [Global Boundary Conditions for Aerodynamics and Aeroacoustic Computations]

Ponnampalam Balakumar - Ph.D., Aeronautics and Astronautics, Massachusetts Institute of Technology, 1986. Associate Professor, Department of Aerospace Engineering, Old Dominion University. Fluid Mechanics [Stability and Transition]

H. Thomas Banks - Ph.D., Applied Mathematics, Purdue University, 1967. Professor, Department of Mathematics, Center for Research in Scientific Computations, North Carolina State University. Applied & Numerical Mathematics [Control Theory]

Richard W. Barnwell - Ph.D., Engineering Mechanics, Virginia Polytechnic Institute and State University, 1968. Professor, Department of Aerospace and Ocean Engineering, Engineering Science and Mechanics, Virginia Polytechnic Institute and State University. Fluid Mechanics [Turbulence Modeling]

Oktay Baysal - Ph.D., Mechanical Engineering, Louisiana State University, 1982. Eminent Scholar and Professor, Department of Aerospace Engineering, Old Dominion University. Applied & Numerical Mathematics

Achi Brandt - Ph.D., Mathematics, The Weizmann Institute of Science, 1965. Professor, Department of Applied Mathematics, The Weizmann Institute of Science, Israel. Applied & Numerical Mathematics [Convergence Acceleration]

Ayodeji O. Demuren - Ph.D., Mechanical Engineering, Imperial College London, United Kingdom, 1979. Associate Professor, Department of Mechanical Engineering and Mechanics, Old Dominion University. Fluid Mechanics [Numerical Modeling of Turbulent Flows]

David Gottlieb - Ph.D., Numerical Analysis, Tel-Aviv University, Israel, 1972. Ford Foundation Professor & Chair, Division of Applied Mathematics, Brown University. Applied & Numerical Mathematics [Boundary Conditions for Hyperbolic Systems]

Chester E. Grosch - Ph.D., Physics and Fluid Dynamics, Stevens Institute of Tchnology, 1967. Professor, Department of Computer Science and Slover Professor, Department of Oceanography, Old Dominion University. Fluid Mechanics [Turbulence and Acoustics]

Jan S. Hesthaven - Ph.D., Applied Mathematics/Numerical Analysis, Technical University of Denmark, 1995. Visiting Assistant Professor, Division of Applied Mathematics, Brown University. Physical Sciences [Computational Electromagnetics]

Fang Q. Hu - Ph.D., Applied Mathematics, Florida State University, 1990. Assistant Professor, Department of Mathematics and Statistics, Old Dominion University. Fluid Mechanics [Aeroacoustics]

Osama A. Kandil - Ph.D., Engineering Mechanics, Virginia Polytechnic Institute and State University, 1974. Professor and Eminent Scholar of Aerospace Engineering and Chair, Department of Aerospace Engineering, Old Dominion University. Applied & Numerical Analysis [Computational Fluid Dynamics]

Frank Kozusko - Ph.D., Computational and Applied Mathematics, Old Dominion University, 1995. Assistant Professor, Department of Mathematics, Hampton University. Fluid Mechanics [Airfoil Design]

David G. Lasseigne - Ph.D., Applied Mathematics, Northwestern University, 1985. Assistant Professor, Department of Mathematics and Statistics, Old Dominion University. Fluid Mechanics [Asymptotic and Numerical Methods for Computational Fluid Dynamics]

Richard W. Longman - Ph.D., Aerospace Engineering, University of California-San Diego, 1969. Professor, Department of Mechanical Engineering, Columbia University. Physical Science [System/Disturbance Identification for Flow Control]

Kurt Maly - Ph.D., Computer Science, Courant Institute, New York University, 1973. Kaufman Professor and Chair, Department of Computer Science, Old Dominion University. Computer Science [High Performance Communication]

James E. Martin - Ph.D., Applied Mathematics, Brown University, 1991. Assistant Professor, Department of Mathematics, Christopher Newport University. Fluid Mechanics [Turbulence and Computation]

Jan Nordstrom - Ph.D., Numerical Analysis, Uppsala University, Sweden, 1993. Senior Scientist, The Aeronautical Research Institute of Sweden. Applied & Numerical Mathematics [Global Boundary Conditions for Aerodynamic and Aeroacoustic Computations]

Alex Pothen - Ph.D., Applied Mathematics, Cornell University, 1984. Professor, Department of Computer Science, Old Dominion University. Computer Science [Parallel Numerical Algorithms]

Viktor Ryaben'kii - Ph.D., Stability of Difference Equations, Moscow State University, 1953. Leading Research Scientist, Keldysh Institute for Applied Mathematics, Russian Academy of Sciences and Full Professor, Department of Control and Applied Mathematics, Moscow Institute of Physics and Technology. Applied & Numerical Mathematics [Global Boundary Conditions for Aerodynamic and Aeroacoustic Computations]

Chi-Wang Shu - Ph.D., Mathematics, University of California-Los Angeles, 1986. Associate Professor, Division of Applied Mathematics, Brown University. Fluid Mechanics [Computational Aeroacoustics]

Ralph C. Smith - Ph.D., Mathematics, Montana State University, 1990. Assistant Professor, Department of Mathematics, Iowa State University. Applied & Numerical Mathematics [Optimal Control Techniques for Structural Acoustics Problems]

Siva Thangam - Ph.D., Mechanical Engineering, Rutgers University, 1980. Professor, Department of Mechanical Engineering, Stevens Institute of Technology. Fluid Mechanics [Turbulence Modeling and Simulation]

Virginia Torczon - Ph.D., Mathematical Sciences, Rice University, 1989. Assistant Professor, Department of Computer Science, The College of William & Mary. Computer Science [Parallel Algorithms for Optimization Including Multidisciplinary Optimization]

Michael W. Trosset - Ph.D., Statistics, University of California-Berkeley, 1983. Department of Mathematics, The College of William & Mary. Applied & Numerical Mathematics [Multidisciplinary Optimization]

Semyon V. Tsynkov - Ph.D., Computational Mathematics, Keldysh Institute for Applied Mathematics, Russian Academy of Sciences, 1991. Senior Lecturer, Department of Applied Mathematics, Tel Aviv University, Israel. Applied & Numerical Mathematics [Global Boundary Conditions for Aerodynamic and Aeroacoustic Computations]

Robert G. Voigt - Ph.D., Mathematics, University of Maryland, 1969. Professor, Computational Science Program, The College of William & Mary. Computer Science [High Performance Computing]

Xiaodong Zhang - Ph.D., Computer Sciences, University of Colorado-Boulder, 1989. Professor, Department of Computer Science, The College of William & Mary. Computer Science

Hans Zima - Ph.D., Mathematics, University of Vienna, Austria, 1964. Professor, Institute for Software Technology and Parallel Systems, University of Vienna, Austria. Computer Science [Compiler Development for Parallel and Distributed Multiprocessors]

Mohammad Zubair - Ph.D., Computer Science, Indian Institute of Technology, Delhi, India, 1987. Professor, Department of Computer Science, Old Dominion University. Computer Science [Performance of Unstructured Flow-solvers on Multi-processor Machines]

X. GRADUATE STUDENTS

Ahmed H. Al-Theneyan - Department of Computer Science, Old Dominion University. (May 1999 to Present)

Abdelkader Baggag - Department of Computer Science, Purdue University. (September 1995 to Present)

Jinghua Fu - Department of Computer Science, Old Dominion University. (June 1999 to Present)

Kerry S. Huston - Department of Mechanical Engineering, Virginia Polytechnic Institute and State University. (May 1999 to August 1999)

David A. Hysom - Department of Computer Science, Old Dominion University. (October 1997 to Present)

Michele Joyner - Department of Mathematics, Center for Research in Scientific Computation, North Carolina State University. (June 3-9, 1999)

Hye-Young Kim - Department of Aerospace Engineering, Texas A&M University. (January 1999 to Present)

Ken Mattsson - Department of Scientific Computing, Uppsala University, Sweden. (May 1999)

Seth D. Milder - Department of Physics and Astronomy, George Mason University. (September 1997 to Present)

Ruben Montero - Departamento de Arquitectura de Computadores y Automatica, Universidad Complutense, Madrid, Spain. (July 1999 to August 1999)

Kara Schumacher Olson - Department of Computer Science, Old Dominion University. (January 1999 to August 2000)

Juan A. Pelaez - Department of Aerospace Engineering, Old Dominion University. (March 1999 to Present)

Manuel Prieto-Matias - Departamento de Arquitectura de Computadores y Automatica, Universidad Complutense, Madrid, Spain. (July 1999 to August 1999)

Kevin Roe - Department of Computer Science, The College of William & Mary. (May 1995 to May 1999)

Radu I. Siminiceanu - Department of Computer Science, The College of William & Mary. (July 1999 to Present)

David Venditti - Department of Aeronautics and Astronautics, Massachusetts Institute of Technology. (June 1999 to July 1999)

Li Xiao - Department of Computer Science, The College of William & Mary. (June 1999 to August 1999)

XI. STUDENT ASSISTANTS

Qian Zhou - School of Engineering and Applied Science, University of Virginia. (June 1999 to July 1999)

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY *(Leave blank)*	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	November 1999	Contractor Report

4. TITLE AND SUBTITLE
Semiannual Report
April 1, 1999 through September 30, 1999

5. FUNDING NUMBERS
C NAS1-97046
WU 505-90-52-01

6. AUTHOR(S)

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Institute for Computer Applications in Science and Engineering
Mail Stop 132C, NASA Langley Research Center
Hampton, VA 23681-2199

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
National Aeronautics and Space Administration
Langley Research Center
Hampton, VA 23681-2199

10. SPONSORING/MONITORING AGENCY REPORT NUMBER
NASA/CR-1999-209728

11. SUPPLEMENTARY NOTES
Langley Technical Monitor: Dennis M. Bushnell
Final Report

12a. DISTRIBUTION/AVAILABILITY STATEMENT
Unclassified–Unlimited
Subject Category 59
Distribution: Nonstandard
Availability: NASA-CASI (301) 621-0390

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*
This report summarizes research conducted at the Institute for Computer Applications in Science and Engineering in applied mathematics, fluid mechanics, and computer science during the period April 1, 1999 through September 30, 1999.

14. SUBJECT TERMS
applied mathematics, multidiscipinary design optimization, fluid mechanics, turbulence, flow control, acoustics, computer science, system software, systems engineering, parallel algorithms

15. NUMBER OF PAGES
69

16. PRICE CODE
A04

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified		